The Art Gallery Problem is $\exists\mathbb{R}$-complete

Mikkel Abrahamsen[1], Anna Adamaszek[1], and Tillmann Miltzow[2]

[1]University of Copenhagen, Denmark. {miab,anad}@di.ku.dk
[2] Université libre de Bruxelles (ULB), Brussels, Belgium. t.miltzow@gmail.com

arXiv:1704.06969v1 [cs.CG] 23 Apr 2017

Abstract

We prove that the *art gallery problem* is equivalent under polynomial time reductions to deciding whether a system of polynomial equations over the real numbers has a solution. The art gallery problem is a classical problem in computational geometry, introduced in 1973 by Viktor Klee. Given a simple polygon $\mathcal{P}$ and an integer k, the goal is to decide if there exists a set G of k *guards* within $\mathcal{P}$ such that every point $p \in \mathcal{P}$ is seen by at least one guard $g \in G$. Each guard corresponds to a point in the polygon $\mathcal{P}$, and we say that a guard g *sees* a point p if the line segment pg is contained in $\mathcal{P}$.

The art gallery problem has stimulated a myriad of research in geometry and in algorithms. However, despite extensive research, the complexity status of the art gallery problem has not been resolved. It has long been known that the problem is NP-hard, but no one has been able to show that it lies in NP. Recently, the computational geometry community became more aware of the complexity class $\exists\mathbb{R}$. The class $\exists\mathbb{R}$ consists of problems that can be reduced in polynomial time to the problem of deciding whether a system of polynomial equations with integer coefficients and any number of real variables has a solution. It can be easily seen that NP $\subseteq \exists\mathbb{R}$. We prove that the art gallery problem is $\exists\mathbb{R}$-complete, implying that (1) any system of polynomial equations over the real numbers can be encoded as an instance of the art gallery problem, and (2) the art gallery problem is not in the complexity class NP unless NP $= \exists\mathbb{R}$. As a corollary of our construction, we prove that for any real algebraic number α there is an instance of the art gallery problem where one of the coordinates of the guards equals α in any guard set of minimum cardinality. That rules out many geometric approaches to the problem.

1 Introduction

We prove that the art gallery problem is equivalent under polynomial time reductions to deciding whether a system of polynomial equations over the real numbers has a solution.

The art gallery problem. Given a simple polygon $\mathcal{P}$, we say that two points $p, q \in \mathcal{P}$ *see each other* if the line segment pq is contained in $\mathcal{P}$. A set of points $G \subseteq \mathcal{P}$ is said to *guard* the polygon $\mathcal{P}$ if every point $p \in \mathcal{P}$ is seen by at least one guard $g \in G$. Such a set G is called a *guard set* of $\mathcal{P}$, and the points of G are called *guards*. A guard set of $\mathcal{P}$ is *optimal* if it is a minimum cardinality guard set of $\mathcal{P}$.

In the *art gallery problem* we are given an integer g and a polygon $\mathcal{P}$ with corners at rational coordinates, and the goal is to decide if $\mathcal{P}$ has a guard set of cardinality g. We consider a polygon as a Jordan curve consisting of finitely many line segments *and* the region that it encloses. The art gallery problem has been introduced in 1973 by Viktor Klee, and it has stimulated a myriad of research in geometry and in algorithms. However, the complexity status of the art gallery problem has stayed unresolved. We are going to prove that the problem is $\exists\mathbb{R}$-complete. Below, we give a formal definition of the complexity class $\exists\mathbb{R}$.

The complexity class $\exists\mathbb{R}$. The *first order theory of the reals* is a set of all true sentences involving real variables, universal and existential quantifiers, boolean and arithmetic operators, constants 0 and 1, parenthesis, equalities and inequalities $(x_1, x_2, \ldots, \forall, \exists, \wedge, \vee, \neg, 0, 1, +, -, \cdot, (\ ,\ )\ , =, <, \leq)$. A formula is called a *sentence* if it has no free variables, i.e., each variable present in the formula is bound by a quantifier. Note that within such formulas one can easily express integer constants (using binary expansion) and powers. Each formula can be converted to a *prenex form*, i.e., a form where it starts with all the quantifiers and is followed by a quantifier-free formula, by a transformation which changes the length of the formula by at most a constant factor.

The *existential theory of the reals* is a set of all true sentences of the first-order theory of the reals in prenex form with existential quantifiers only, i.e., sentences of the form

$$(\exists X_1 \exists X_2 \ldots \exists X_k)\ \Phi(X_1, X_2, \ldots, X_k),$$

where Φ is a quantifier-free formula of the first-order theory of the reals with variables $X_1, \ldots, X_k$. The problem ETR is the problem of deciding whether a given existential formula of the above form is true. The complexity class $\exists\mathbb{R}$ consists of all problems that are reducible to ETR in polynomial time. The most well-known problem in the complexity class $\exists\mathbb{R}$ is deciding whether a system of polynomial equations over the real numbers has a solution.

It is currently known that

$$\text{NP} \subseteq \exists\mathbb{R} \subseteq \text{PSPACE}.$$

It is not hard see that the problem ETR is NP-hard, for instance by the following reduction from 3SAT. For each boolean variable x in an instance of 3SAT, we introduce a real variable x', and require that $x' \cdot (1 - x') = 0$ in order to ensure that $x' \in \{0, 1\}$. For any clause of the 3SAT formula we construct a function which evaluates to 1 if the corresponding clause is satisfied, and to 0 otherwise. For a clause C of the form $x \vee y \vee \neg z$, the corresponding function C' is $1 - (1 - x')(1 - y')z'$. The conjunction of clauses $C_1 \wedge \ldots \wedge C_m$ is then translated to the equation $C'_1 \cdot \ldots \cdot C'_m - 1 = 0$. Clearly, a formula of 3SAT is true if and only if the constructed set of equations has a solution in $\mathbb{R}$. The containment $\exists\mathbb{R} \subseteq$ PSPACE is highly non-trivial, and it has first been established by Canny [10].

By the reduction from 3SAT to ETR sketched above we know that a problem of deciding whether a given polynomial equation over $\{0, 1\}$ with integer coefficients has a solution is NP-hard. The problem is also in NP, as a satisfying assignment clearly serves as a witness.

Therefore, NP-complete problems are the problems equivalent (under polynomial time reductions) to deciding whether a given polynomial equation over $\{0, 1\}$ with integer coefficients has a solution. A well-known $\exists\mathbb{R}$-complete problem is the problem of deciding whether a single polynomial equation $Q(x_1, \ldots, x_n) = 0$ with integer coefficients has a solution in $\mathbb{R}$ [27, Proposition 3.2]. Therefore, the $\exists\mathbb{R}$-complete problems are equivalent to deciding whether a given polynomial equation over $\mathbb{R}$ with integer coefficients has a solution.

Our results and their implications. We prove that solving the art gallery problem is, up to a polynomial time reduction, as hard as deciding whether a system of polynomial equations and inequalities over the real numbers has a solution.

Theorem 1. *The art gallery problem is $\exists\mathbb{R}$-complete, even the restricted variant where we are given a polygon with vertices at integer coordinates.*

In our construction, an ETR formula $(\exists X_1 \ldots \exists X_k)\ \Phi(X_1, \ldots, X_k)$ is transformed into an instance $(\mathcal{P}, g)$ of the art gallery problem where $g > k$. Let S_Φ denote the solution space of the formula Φ, i.e., $S_\Phi := \{x \in \mathbb{R}^k : \Phi(x)\}$. We will prove the following theorem.

Theorem 2. *Let Φ be an ETR formula with k variables. Then there is an instance $(\mathcal{P}, g)$ of the art gallery problem, and constants $c_1, d_1, \ldots, c_k, d_k \in \mathbb{Q}$, such that*

- *if Φ has a solution, then $\mathcal{P}$ has a guard set of size g, and*
- *for any guard set G of $\mathcal{P}$ of size g, there exists $(x_1, \ldots, x_k) \in S_\Phi$ such that G contains guards at positions $(c_1x_1 + d_1, 0), \ldots, (c_kx_k + d_k, 0)$.*

We get the following corollary.

Corollary 3. *Given any real algebraic number α, there exists a polygon $\mathcal{P}$ with vertices at rational coordinates such that in any optimal guard set of $\mathcal{P}$ there is a guard with an x-coordinate equal to α.*

It is a classical result in Galois theory, and has thus been known since the 19th century, that there are polynomial equations of degree five with integer coefficients which have real solutions, but with no solutions expressible by radicals (i.e., solutions that can be expressed using integers, addition, subtraction, multiplication, division, raising to integer powers, and the extraction of n'th roots). One such example is the equation $x^5 - x + 1 = 0$ [34]. It is a peculiar fact that using the reduction described in this paper, we are able to transform such an equation into an instance of the art gallery problem where no optimal guard set can be expressed by radicals.

Our result rules out many approaches to solving the art gallery problem. A natural approach to finding a guard set for a given polygon $\mathcal{P}$ is to create a candidate set for the guards, and select a guard set as a subset of the candidate set. For instance, a candidate set can consist of the vertices of $\mathcal{P}$. The candidate set can then be expanded by considering all lines containing two candidates and adding all intersection points of these lines to the candidate set. This process can be repeated any finite number of times, but only candidates with rational coordinates can be obtained that way, and the candidate set will thus not contain an optimal guard set in general. Algorithms of this kind are for instance discussed by de Rezende *et al.* [14]. One can get a more refined set of candidates by also considering certain quadratic curves [5], or more complicated curves. Our results imply that if the algebraic degree of the considered curves is bounded by a constant, a such approach cannot lead to an optimal solution in general, since the coordinates of the candidates will also have algebraic degree bounded by a constant.

Related work. The art gallery problem has been extensively studied, with some books, surveys, and book chapters dedicated to it [30, 40, 44, 32, 13, 26, 15, 31]. The research is stimulated by a large number of possible variants of the problem and related questions that can be studied. The version of the art gallery problem considered in this paper is the classical one, and it has been originally formulated by Victor Klee (see O'Rourke [30]). Other versions of the art gallery problem include restrictions on the positions of the guards, different definitions of visibility, restricted classes of polygons, restricting the part of the polygon that has to be guarded, etc.

The art gallery problem has been studied both from combinatorial and from the algorithmic perspective. Studies have been made on algorithms performing well in practice on real-world and simulated instances of the problem [8, 14]. Another branch of research investigates approximation algorithms for the art gallery problem and its variants [17, 7, 21].

The first exact algorithm for solving the art gallery problem was published in 2002 in the conference version of a paper by Efrat and Har-Peled [16]. They attribute the result to Micha Sharir. Before this time, the problem was not even known to be decidable. The approach is to reduce the art gallery problem to a formula in the first order theory of the reals and use standard algebraic methods to decide if that formula is true, such as the techniques provided by Basu *et al.* [4]. No algorithm is known that avoids the use of this powerful machinery.

Lee and Lin [24] proved, by constructing a reduction from 3SAT, that the art gallery problem is NP-hard when the guards are restricted to the vertices of the polygon. It has subsequently been shown by Aggarwal ([3], see also [30]) that this argument can be extended to the case with no restrictions on the guards. Various papers showed other hardness results or conditional lower bounds for the art gallery problem and its variations [24, 39, 43, 9, 17, 33, 23, 7, 20].

Note that the version of the art gallery problem where the guards are restricted to the vertices of the polygon is obviously in NP. Another related problem is the *terrain guarding problem*. Here, the area above an x-monotone polygonal curve c has to be guarded by guards restricted to c. Friedrichs *et al.* [18] recently showed that terrain guarding is in NP. These two problems are both NP-complete. The authors of the present paper [2] gave a simple example of a polygon with a unique optimal guard set consisting of three guards at irrational coordinates. Four guards are needed if they have to have rational coordinates. This could be an indication that the original version of the art gallery problem with no restrictions on the guards is actually more difficult than these related problems. Friedrichs *et al.* [18] stated that "[...] it is a long-standing open problem for the more general Art Gallery Problem (AGP): For the AGP it is not known whether the coordinates of an optimal guard cover can be represented with a polynomial number of bits". In the present paper, we show that such a representation does not exist and that the art gallery problem is indeed not in NP under the assumption NP $\neq \exists\mathbb{R}$.

A growing class of problems turn out to be equivalent (under polynomial time reductions) to deciding whether polynomial equations and inequalities over the reals have a solution. These problems form the family of $\exists\mathbb{R}$-complete problems as it is currently known. This class includes problems like the stretchability of pseudoline arrangements [29, 42], recognition of intersection graphs of various objects (e.g. segments [27], unit disks [28], and general convex sets [36]), recognition of point visibility graphs [12], the Steinitz problem for 4-polytopes [35], deciding whether a graph with given edge lengths can be realized by a straight-line drawing [37, 1], deciding whether a straight line drawing of a graph exists with a given number of edge crossings [6], decision problems related to Nash-equilibria [19], and positive semidefinite matrix factorization [41]. We refer the reader to the lecture notes by Matoušek [27] and surveys by Schaefer [36] and Cardinal [11] for more information on the complexity class $\exists\mathbb{R}$.

Overview of the paper and techniques. In Section 2 we show that the art gallery problem is in the complexity class $\exists\mathbb{R}$. For that we present a construction of an ETR-formula Φ for any instance $(\mathcal{P}, g)$ of the art gallery problem such that Φ has a solution if and only if $\mathcal{P}$ has a

guard set of size g. The details of the construction are in Appendix A. The idea is to encode guards by pairs of variables and compute a set of witnesses (which depend on the positions of the guards) of polynomial size such that the polygon is guarded if and only if the witnesses are seen by the guards.

The proof that the art gallery problem is $\exists\mathbb{R}$-hard is the main result of the paper, and it consists of two parts. The first part is of algebraic nature, and in that we introduce a novel $\exists\mathbb{R}$-complete problem which we call ETR-INV. A common way of making a reduction from ETR to some other problem is to build gadgets corresponding to each of the equations $x = 1$, $x + y = z$, and $x \cdot y = z$ for any variables x, y, z. Usually, the multiplication gadget is the most involved one. An instance of ETR-INV is a conjunction of formulas of the form $x = 1$, $x + y = z$, and $x \cdot y = 1$, with the requirement that each variable must be in the interval $[1/2, 2]$. In particular, the reduction from ETR-INV requires building a gadget for *inversion* (i.e., $x \cdot y = 1$) instead of a more general gadget for multiplication. The formal definition of ETR-INV and the proof that it is $\exists\mathbb{R}$-complete is presented in Section 3 (with details in Appendix B). We think that the problem ETR-INV might be of independent interest, and that it will simplify constructing $\exists\mathbb{R}$-hardness proofs. The problem ETR-INV has already been used to prove $\exists\mathbb{R}$-completeness of a geometric graph drawing problem with prescribed face areas [22].

In Section 4 (with details in Appendix C) we describe a polynomial time reduction from ETR-INV to the art gallery problem, which shows that the art gallery problem is $\exists\mathbb{R}$-hard. This reduction constructs an art gallery instance $(\mathcal{P}(\Phi), g(\Phi))$ from an ETR-INV instance Φ, such that $\mathcal{P}(\Phi)$ has a guard set of size $g(\Phi)$ if and only if the formula Φ has a solution. We construct the polygon so that it contains $g(\Phi)$ *guard segments* (which are horizontal line segments within $\mathcal{P}$) and *stationary guard positions* (points within $\mathcal{P}$). By introducing *pockets* we enforce that if $\mathcal{P}$ has a guard set of size $g(\Phi)$, then there must be exactly one guard at each guard segment and at each stationary guard position. Each guard segment represents a variable of Φ (with multiple segments representing the same variable) in the sense that the position of the guard on the segment specifies the value of the variable, the endpoints of a segment corresponding to the values $1/2$ and 2.

We develop a technique for *copying* guard segments, i.e., enforcing that the guards at two segments correspond to the same variable. We do that by introducing *critical segments* within the polygon, which can be seen by guards from two guard segments (but not from other guard segments). Then the requirement that a critical segment is seen introduces dependency between the guards at the corresponding segments. Different critical segments will enforce different dependencies, and by enforcing that two guards must see together two particular critical segments, we can ensure that the guards represent the same value. The stationary guards are placed to see the remaining areas of the polygon.

With this technique, we are able to copy two or three segments from an area containing guard segments corresponding to all variables into a *gadget*, where we will enforce a dependency between the values of the variables represented by the two or three segments. This is done by constructing a *corridor* containing two critical segments for each pair of copied segments. The construction is non-trivial, as it requires the critical segments not to be seen from any other segments.

Within the gadgets, we build features that enforce the variables x, y, z represented by the guards to satisfy one of the conditions $x + y \geq z$, $x + y \leq z$, or $x \cdot y = 1$. The conditions are enforced by a requirement that two or three guards can see together some areas, where for the case of a gadget with three variables the area to be seen is a quadrilateral instead of a line segment.

2 The art gallery problem is in $\exists\mathbb{R}$

In this section we sketch a proof that the art gallery problem is in the complexity class $\exists\mathbb{R}$. Our proof also works for the more general version of the art gallery problem where the input polygon can have polygonal holes.

Theorem 4. *The art gallery problem is in the complexity class* $\exists\mathbb{R}$.

Sketch of proof. Let $(\mathcal{P}, g)$ be an instance of the art gallery problem where the polygon $\mathcal{P}$ has n vertices, each of which has rational coordinates represented by at most B bits. Assume that $g < n$, as it is otherwise trivial that $\mathcal{P}$ can be guarded by g guards. We show how to construct in polynomial time a quantifier-free formula $\Phi := \Phi(\mathcal{P}, g)$ of the first-order theory of the reals such that Φ is satisfiable if and only if $\mathcal{P}$ has a guard set of cardinality g. It has been our priority to define the formula Φ so that it is as simple as possible to describe. It might be possible to construct an equivalent but shorter formula.

The description of Φ resembles the formula Ψ that Micha Sharir described to Efrat and Har-Peled [16]:

$$\Psi \;:=\; \left[\exists x_1, y_1, \ldots, x_g, y_g \; \forall p_x, p_y : \text{INSIDE-POLYGON}(p_x, p_y) \Longrightarrow \bigvee_{i=1}^{g} \text{SEES}(x_i, y_i, p_x, p_y)\right].$$

For each $i \in \{1, \ldots, g\}$, the variables x_i, y_i represent the position of a guard $g_i := (x_i, y_i)$, and $p := (p_x, p_y)$ represents an arbitrary point. The predicate INSIDE-POLYGON(p_x, p_y) tests if the point p is contained in the polygon $\mathcal{P}$ and SEES(x_i, y_i, p_x, p_y) checks if the guard g_i can see the point p. Thus, the formula is satisfiable if and only if $\mathcal{P}$ has a guard set of cardinality g. Note that although the implication "$\Longrightarrow$" is not allowed in the first order theory of the reals, we can always substitute "$A \Longrightarrow B$" by "$\neg A \vee B$".

The formula Ψ is not an existential formula. Our main idea behind obtaining a formula with no universal quantifier is finding a polynomial number of points inside $\mathcal{P}$ which, if all seen by the guards, ensure that all of $\mathcal{P}$ is seen. We denote such a set of points as a *witness set*.

Creating a witness set. We are now ready to describe the witness set that replaces the universal quantifier. Let $\mathcal{L} := \{\ell_1, \ldots, \ell_m\}$ be the set of lines containing either an edge of $\mathcal{P}$, or a guard g_i for $i \in \{1, \ldots, g\}$ and a vertex v of $\mathcal{P}$.* The well-defined lines in $\mathcal{L}$ partition the plane into a collection of open *regions* $\mathcal{A}$, which are connected components of $\mathbb{R}^2 \setminus \bigcup_{\ell \in \mathcal{L}} \ell$ (see Figure 1 for an example).

The set $\mathcal{A}$ has the following properties.

- Each region in $\mathcal{A}$ is an open convex polygon.
- Each region in $\mathcal{A}$ is either contained in $\mathcal{P}$ or contained in the complement of $\mathcal{P}$.
- The closure of the union of the regions that are contained in $\mathcal{P}$ equals $\mathcal{P}$.
- For each region $R \in \mathcal{A}$, each guard g_i either sees all points of R, or sees no point of R. In particular, if a guard g_i sees one point in R, it sees all of R and its closure.

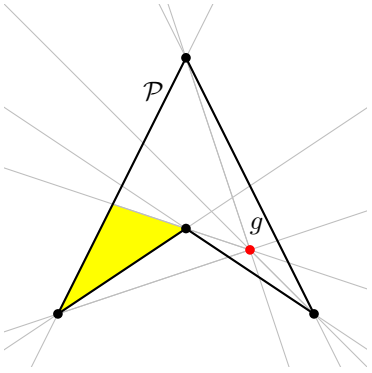

Figure 1: A polygon $\mathcal{P}$ together with a guard g and the resulting line arrangement $\mathcal{L}$. The yellow region is not seen.

*Note that if a line is defined as passing through a guard g_i and a vertex v such that g_i and v are coincident, the line is not well-defined. Such lines are not considered in the partition into regions described below, but they are included in the set $\mathcal{L}$. However, in this sketch of the proof, we will ignore them.

Thus, it is sufficient to test that for each region $R \in \mathcal{A}$ which is in $\mathcal{P}$, at least one point in R is seen by a guard. For three points a, b, c, define the *centroid* of a, b, c to be the point $C(a, b, c) := (a + b + c)/3$. If a, b, c are three different vertices of the same region $R \in \mathcal{A}$, then the centroid $C(a, b, c)$ must lie in the interior of R. Note that each region has at least three vertices and thus contains at least one such centroid. Let X be the set of all intersection points of two non-parallel, well-defined lines in $\mathcal{L}$, i.e., X consists of all vertices of all regions in $\mathcal{A}$. In the formula Φ, we generate all points in X. For any three points a, b, c in X, we also generate the centroid $C(a, b, c)$. If the centroid is in $\mathcal{P}$, we check that it is seen by a guard. Since there are $O(((gn)^2)^3)$ centroids of three points in X and each is tested by a formula of size $O(gnB)$, we get a formula of length $O(g^7 n^7 B^2) = O(n^{14} B^2)$. Clearly, the formula can be computed in polynomial time. □

3 The problem ETR-INV

To show that the art gallery problem is $\exists\mathbb{R}$-hard, we will provide a reduction from the problem ETR-INV, which we introduce below. In this section, we sketch how to show that ETR-INV is $\exists\mathbb{R}$-complete.

Definition 5 (ETR-INV)**.** *In the problem ETR-INV, we are given a set of real variables* $\{x_1, \ldots, x_n\}$*, and a set of equations of the form*

$$x = 1, \quad x + y = z, \quad x \cdot y = 1,$$

for $x, y, z \in \{x_1, \ldots, x_n\}$*. The goal is to decide whether the system of equations has a solution when each variable is restricted to the range* $[1/2, 2]$*.*

Theorem 6. *The problem ETR-INV is* $\exists\mathbb{R}$*-complete.*

Sketch of proof. To show that ETR-INV is $\exists\mathbb{R}$-hard, we perform a series of polynomial time reductions, starting from a formula $(\exists X_1 \ldots \exists X_k)\ \Phi(X_1, \ldots, X_k)$ which is an instance of the problem ETR. We use a lemma by Schaefer and Štefankovič [38] to reduce this formula to a single polynomial F with integer coefficients, so that the equation $F = 0$ has a solution if and only if Φ has a solution. Using a standard technique, we transform the equation $F = 0$ into a formula Φ' which is a conjunction of equations of the forms $x = 1$, $x + y = z$, and $x \cdot y = z$. The idea is to replace each addition or multiplication of two variables x, y in F by a new variable z, and to add to Φ' the corresponding condition $x + y = z$ or $x \cdot y = z$. The coefficients of F are decomposed in a similar way. Another standard technique is used to get an equivalent problem where the variables have been scaled down to the range $[-1/8, 1/8]$. A third standard technique is used to shift the variables to the range $[1/2, 2]$. We finally show how we can substitute each equation of the form $x \cdot y = z$ by an equivalent set of equations using only addition and inversion, i.e., equations of the forms $x + y = z$ and $x \cdot y = 1$. This last step is our novel contribution in the proof of $\exists\mathbb{R}$-hardness of ETR-INV, and we explain the main ideas in the following. In this sketch, we ignore that the variables should stay in the range $[1/2, 2]$.

We first show how to define a new variable V_{x^2} satisfying $V_{x^2} = x^2$. The technique is based on the observation that for $x \notin \{0, 1\}$, we have $\frac{1}{x-1} - \frac{1}{x} = \frac{1}{x^2-x}$. With the equations allowed in ETR-INV we can easily construct variables $V_{\frac{1}{x-1}}$ and $V_{\frac{1}{x}}$ satisfying $V_{\frac{1}{x-1}} = \frac{1}{x-1}$ and $V_{\frac{1}{x}} = \frac{1}{x}$. The equation $V_{\frac{1}{x}} + V_{\frac{1}{x^2-x}} = V_{\frac{1}{x-1}}$ ensures that for a new variable $V_{\frac{1}{x^2-x}}$, we have $V_{\frac{1}{x^2-x}} = \frac{1}{x^2-x}$. Now, the equations $V_{x^2-x} \cdot V_{\frac{1}{x^2-x}} = 1$ and $V_{x^2-x} + x = V_{x^2}$ ensure that $V_{x^2} = x^2$.

In order to substitute equations of the form $x \cdot y = z$ with the equations allowed in ETR-INV, we observe that $(x + y)^2 - x^2 - y^2 = 2xy$. We can create variables satisfying $V_{(x+y)^2} = (x + y)^2$, $V_{x^2} = x^2$, and $V_{y^2} = y^2$ using the technique for squaring a variable described before. It is then easy to construct a variable V_{2xy} satisfying $V_{2xy} = 2xy$.

The constructed instance of ETR-INV can be computed in polynomial time, and it has a solution if and only if the formula Φ has a solution. Therefore, ETR-INV is $\exists\mathbb{R}$-hard. As the conjunction of the equations of ETR-INV, together with the inequalities describing the restricted range $[1/2, 2]$ of the variables, is a quantifier-free formula of the first-order theory of the reals, ETR-INV is in $\exists\mathbb{R}$, which yields that ETR-INV is $\exists\mathbb{R}$-complete. □

4 Reduction from ETR-INV to the art gallery problem

Overview of the construction. Let Φ be an instance of ETR-INV consisting of k equations with n variables $X := \{x_0, \ldots, x_{n-1}\}$. We show that there exists a polygon $\mathcal{P} := \mathcal{P}(\Phi)$ with corners at rational coordinates which can be computed in polynomial time such that Φ has a solution if and only if $\mathcal{P}$ can be guarded by some number $g := g(\Phi)$ of guards. The number g will follow from the construction. A sketch of the polygon $\mathcal{P}$ is shown in Figure 2.

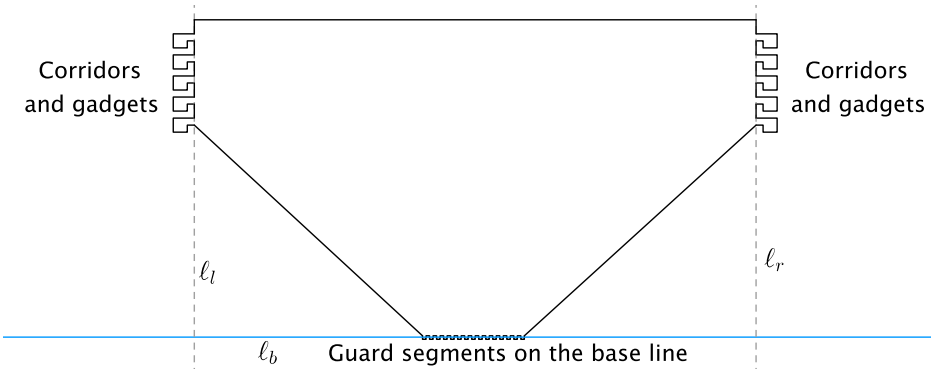

Figure 2: A high-level sketch of the construction of the polygon $\mathcal{P}$.

Each variable $x_i \in X$ is represented by a collection of *guard segments*, which are horizontal line segments contained in the interior of $\mathcal{P}$. Consider one guard segment $s := ab$, where a is to the left of b, and assume that s represents the variable x_i and that there is exactly one guard p placed on s. The guard segment s can be oriented to the right or to the left. The guard p on s specifies the value of the variable x_i as $\frac{1}{2} + \frac{3\|ap\|}{2\|ab\|}$ if s is oriented to the right, and $\frac{1}{2} + \frac{3\|bp\|}{2\|ab\|}$ if s is oriented to the left. Here, the additive term $\frac{1}{2}$ and the factor $\frac{3}{2}$ stem from the fact that all the variables in X are contained in the interval $[\frac{1}{2}, 2]$.

Suppose that there exists a solution to Φ. We will show that in that case any optimal guard set G of $\mathcal{P}$ has size $g(\Phi)$ and *specifies* a solution to Φ in the sense that it satisfies the following two properties.

- Each variable $x_i \in X$ is specified *consistently* by G, meaning that there is exactly one guard on each guard segment representing x_i, and all these guards specify the same value of x_i.
- The guard set G is *feasible*, meaning that the values of X thus specified is a solution to Φ.

Moreover, if there is no solution to Φ, each guard set of $\mathcal{P}$ consists of more than $g(\Phi)$ guards.

The polygon $\mathcal{P}$ is constructed in the following way. The bottom part of the polygon consists of a collection of *pockets*, containing in total $4n$ collinear and equidistant guard segments $s_0, \ldots, s_{4n-1}$, three right-oriented and one left-oriented segment corresponding to each variable of Φ. We denote the horizontal line containing these guard segments as the *base line* or ℓ_b. At the left and at the right side of $\mathcal{P}$, there are some *corridors* attached, each of which leads into a *gadget*. The *entrances* to the corridors at the right side of $\mathcal{P}$ are line segments contained in a vertical line ℓ_r. Likewise, the entrances to the corridors at the left side of $\mathcal{P}$ are contained

in a vertical line ℓ_l. The gadgets also contain guard segments, and they are used to impose dependencies between the guards in order to ensure that if there is a solution to Φ, then any minimum guard set of $\mathcal{P}$ consists of $g(\Phi)$ guards and specifies a solution to Φ in the sense defined above. The corridors are used to copy the positions of guards on guard segments on the base line to guards on guard segments inside the gadgets. Each gadget corresponds to a constraint of one of the types $x + y \geq z$, $x + y \leq z$, $x \cdot y = 1$, $x + y \geq 5/2$, and $x + y \leq 5/2$. The first three types of constraints are used to encode the dependencies between the variables in X as specified by Φ, whereas the latter two constraints are used to encode the dependencies between the right-oriented and left-oriented guard segments representing a single variable in X. The constraints of type $x = 1$ are enforced by modifying the pocket containing a guard segment corresponding to x.

Creating stationary guard positions and guard segments. We denote some points in $\mathcal{P}$ as *stationary guard positions*. A guard placed at a stationary guard position is called a *stationary guard*. A stationary guard position is the unique point $p \in \mathcal{P}$ such that a guard placed at p can see some set of vertices (usually two vertices) of $\mathcal{P}$. We use stationary guards for the purpose of seeing some region on one side of a line segment ℓ, but no points on the other side of ℓ. See Figure 3 (left) for an explanation of how such a construction can be made.

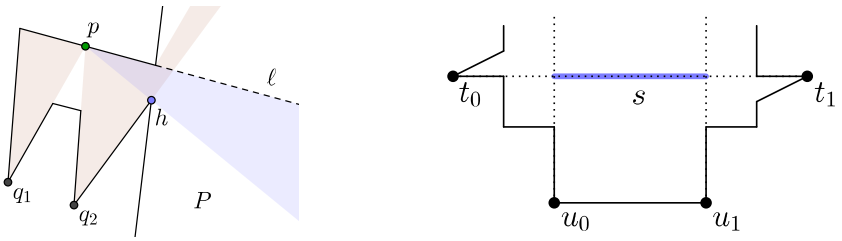

Figure 3: Left: The construction of a stationary guard position p that sees an area in P below a line segment ℓ. The brown areas are the regions of points that see q_1 and q_2, and p is the only point that sees both q_1 and q_2. The point p sees the points in the blue wedge, and the angle of the wedge can be adjusted by choosing the point h accordingly. Right: The construction of a guard segment s (the blue segment). In order to see the points t_0, t_1, a guard must be on the horizontal dotted segment. Furthermore, in order to see u_0, u_1, the guard must be between the vertical dotted segments that contain the endpoints of s. Thus, a guard sees t_0, t_1, u_0, u_1 if and only if the guard is at s.

We likewise denote some horizontal line segments of $\mathcal{P}$ as *guard segments*. A guard segment s consists of all points from which a guard can see some set of four vertices of $\mathcal{P}$. See Figure 3 (right) for an example of such a construction.

In the full version of the reduction, we will prove that for any guard set of size of at most $g(\Phi)$, there is one guard placed at each stationary guard position and at each guard segment, and there are no guards except of these positions. As explained earlier, guards placed at the guard segments will be used to encode the values of the variables of Φ.

Imposing inequalities by nooks and umbras. The nooks and umbras, which we introduce below, are our basic tools used to impose dependency between guards placed on two different guard segments. For the following definitions, see Figure 4.

Definition 7 (nook and umbra)**.** *Let $\mathcal{P}$ be a polygon with guard segments $r_0 := a_0b_0$ and $r_1 := a_1b_1$, where r_0 is to the left of r_1. Let c_0, c_1 be two vertices of $\mathcal{P}$, such that c_0 is to the left of c_1. Suppose that the rays $\overrightarrow{b_0c_0}$ and $\overrightarrow{b_1c_1}$ intersect at a point f_0, the lines $\overrightarrow{a_0c_0}$ and $\overrightarrow{a_1c_1}$ intersect at a point f_1, and that $Q := c_0c_1f_1f_0$ is a convex quadrilateral contained in $\mathcal{P}$. For each $i \in \{0, 1\}$ define the function $\pi_i \colon r_i \longrightarrow f_0f_1$ such that $\pi_i(p)$ is the intersection of the ray $\overrightarrow{pc_i}$ with the line segment f_0f_1, and suppose that π_i is bijective.*

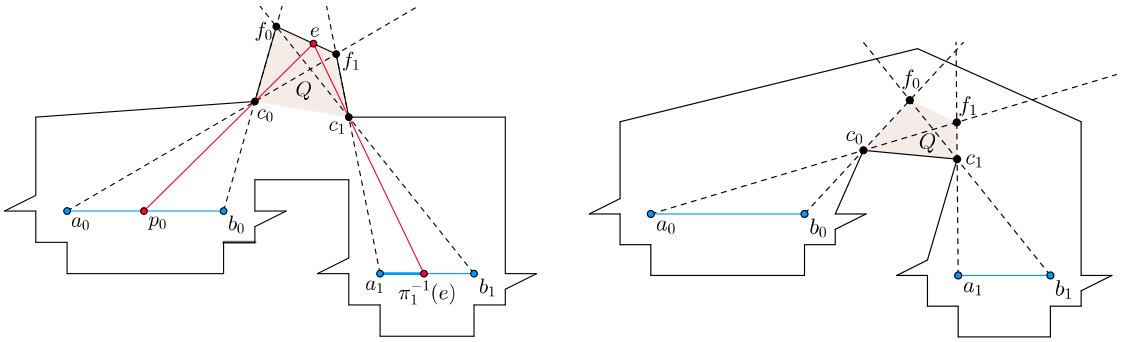

Figure 4: The brown area Q representing a nook (left), and an umbra (right). In the left figure, note that if a guard p_1 placed at the segment a_1b_1 has to see the whole line segment f_0f_1 together with p_0, then p_1 must be on or to the left of the point $\pi_1^{-1}(e)$, where $e := \pi_0(p_0)$.

We say that Q is a nook *of the pair of guard segments r_0, r_1 if for each $i \in \{0,1\}$ and every $p \in r_i$, a guard at p can see all of the segment $\pi_i(p)f_{1-i}$ but nothing else of f_0f_1. We say that Q is an* umbra *of the segments r_0, r_1 if for each $i \in \{0,1\}$ and every $p \in r_i$, a guard at p can see all of the segment $\pi_i(p)f_i$ but nothing else of f_0f_1. The functions π_0, π_1 are called* projections *of the nook or the umbra.*

Definition 8 (critical segment and shadow corners)**.** *Consider a nook or an umbra $Q := c_0c_1f_1f_0$ of a pair of guard segments r_0, r_1. The line segment f_0f_1 is called the* critical segment *of Q, and the vertices c_0, c_1 are called the* shadow corners *of Q.*

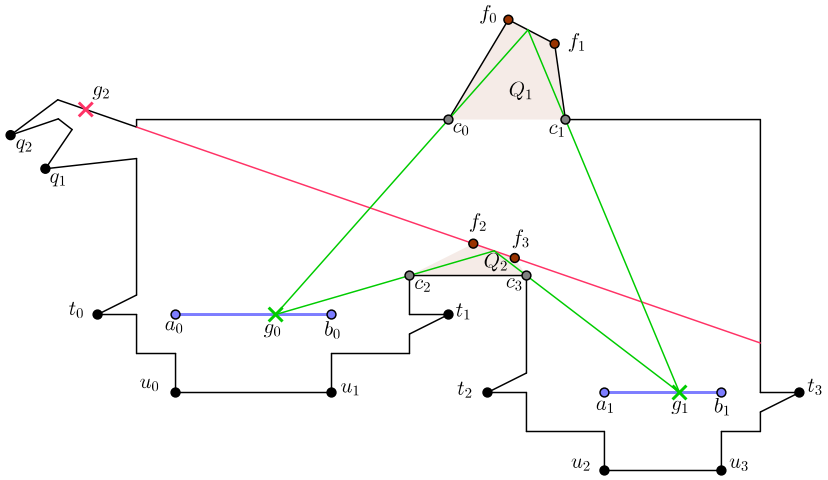

Figure 5: Q_1 is a copy-nook of the segments $r_0 := a_0b_0$ and $r_1 := a_1b_1$ with a critical segment f_0f_1, and Q_2 is a copy-umbra for the same pair with a critical segment f_2f_3. It can be seen that this polygon cannot be guarded by fewer than 3 guards, and any guard set with 3 guards must contain a guard g_0 on r_0, a guard g_1 on r_1, and a stationary guard at the point g_2. The guards g_0 and g_1 must specify the same value on r_0 and r_1, respectively.

We will construct nooks and umbras for pairs of guard segments where we want to enforce dependency between the values of the corresponding variables. When making use of an umbra, we will also create a stationary guard position from which a guard sees the whole quadrilateral Q, but nothing on the other side of the critical segment f_0f_1. In this way we can enforce the guards on r_0 and r_1 to see all of f_0f_1 together. For the case of a nook, the segment f_0f_1 will always be on the polygon boundary, and then there will be no stationary guard needed. See Figure 5 for an example of a construction of both a nook and an umbra for a pair of guard segments.

Definition 9. *Let Q be a nook or an umbra of a pair of guard segments $r_0 := a_0b_0$ and $r_1 := a_1b_1$ with the same orientation, such that the shadow corners c_0 and c_1 have the same y-coordinate. We then call Q a* copy-nook *or a* copy-umbra*, respectively.*

Consider a pair of guard segments r_0, r_1 oriented in the same way, for which there is both a copy-nook and a copy-umbra. We can show that if there is one guard placed at each segment, and the guards together see both critical segments, then the two guards specify the same value. We use this observation to make one guard segment a *copy* of another. See Figure 5, where the guard segment a_1b_1 is a copy of a_0b_0.

Corridors. Inside each gadget there are two or three guard segments r_i, r_j, r_l (or r_i, r_j) corresponding to two or three pairwise different guard segments from the base line s_i, s_j, s_l (or s_i, s_j). A corridor ensures that the segments r_i, r_j, r_l are copies of the segments s_i, s_j, s_l, respectively. To obtain that, we need to construct a copy-nook and a copy-umbra within the corridor for each pair of corresponding segments, see Figure 6 for a simplified illustration. The lower wall of the corridor of the gadget is a horizontal edge c_0, c_1. The vertices c_0, c_1 of the corridor act as shadow corners in three overlapping copy-umbras for the pairs (s_i, r_i), (s_j, r_j), and (s_l, r_l), respectively. We construct the top wall of the corridor so that it creates three copy-nooks for the same pairs. To enforce that for any guard set of size $g(\Phi)$ and each $\sigma \in \{i, j, l\}$, the guard segments s_σ and r_σ specify the same value, we have to ensure that no guards on guard segments other than s_σ and r_σ can see the critical segments of the copy-umbra and the copy-nook of the pair s_σ, r_σ. The precise construction of the corridor that ensures this property is non-trivial. The high-level idea behind the construction is as follows. By placing the corridor sufficiently far away from the segments on the base line, and by making the corridor *entrance* sufficiently small, we obtain that the visibility lines from the guard segment endpoints through the points at the corridor entrance are almost parallel and can be described by a simple pattern. Moreover, each point in the corridor which is far enough from the corridor entrances can be seen by points from at most one guard segment placed to the left of the left entrance, and from at most one guard segment placed to the right of the right entrance. Stationary guards within the corridor ensure that the remaining area of the corridor is seen. This allows us to construct the corridor with the desired properties.

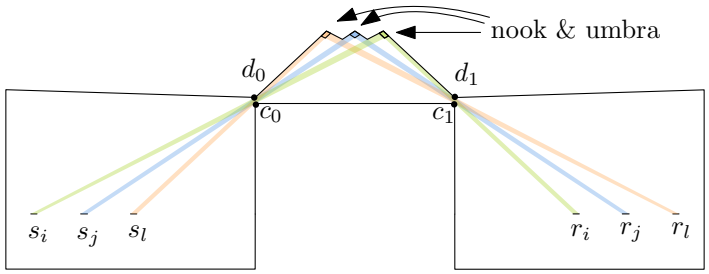

Figure 6: In this figure, we display a simplified corridor construction. The vertices c_0, c_1 serve as shadow corners for three copy-umbras simultaneously for the pairs $(s_i, r_i), (s_j, r_j), (s_l, r_l)$. Each of these pairs also have a small copy-nook in the top of corridor. The entrances c_0d_0 and c_1d_1 to the corridor are sufficiently small so that the critical segments of the nook and umbra of each pair of segments s_σ, r_σ (contained in the small boxes at the top of the figure) are not seen by other guard segments.

Addition gadget. For any equation of the form $x_i + x_j = x_l$ in Φ, where $i, j, l \in \{0, \ldots, n-1\}$, we construct a $\geq$-addition gadget which represents the inequality $x_i + x_j \geq x_l$, and a $\leq$-addition gadget for the inequality $x_i + x_j \leq x_l$.

The general idea behind the construction of the gadget imposing the $\geq$-inequality is as follows (see Figure 7). We can place three guard segments r_i', r_j, r_l in such a way that a certain

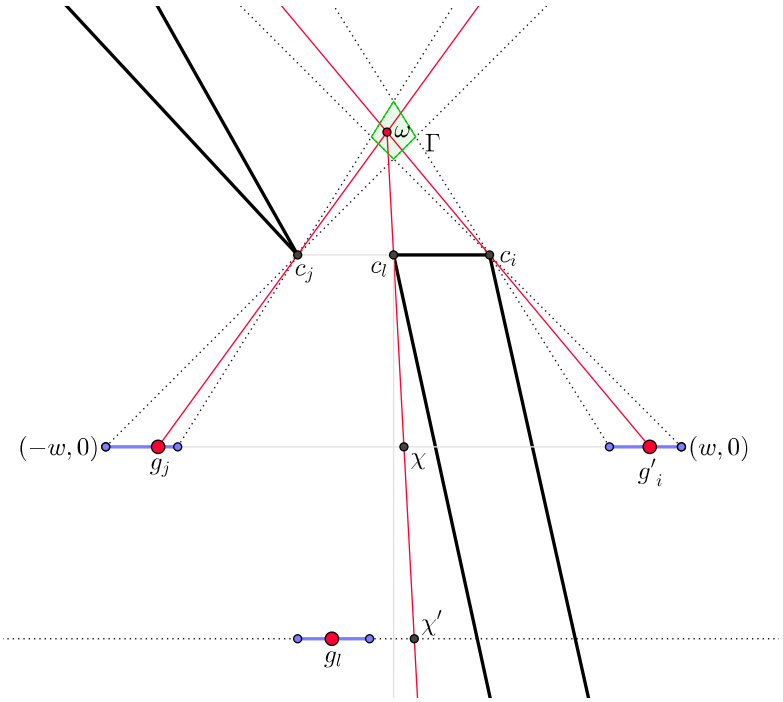

Figure 7: The thick black segments are edges of the polygon. The green quadrilateral is Γ. In order to see Γ together with g'_i and g_j, the guard g_l must be on or to the left of the point χ'.

quadrilateral Γ is seen by the guards g'_i, g_j, g_l placed at r'_i, r_j, r_l, respectively, if and only if the values x'_i, x_j, x_l specified by the corresponding segments satisfy the inequality $x'_i + x_j \geq x_l$. The area of the gadget except of Γ can be seen with the help of stationary guards. The actual gadget is more complicated, as the guard segments r'_i, r_j, r_l need to be copies (where copying is performed via corridor, as described earlier) of the base line guard segments. To make it possible, we introduce an additional guard segment r_i within the gadget, and we copy appropriately chosen segments from the base line into r_i, r_j, r_l from the left (i.e., the gadget is placed in the right side of $\mathcal{P}$). Then, by introducing a copy-nook within the construction sketched in Figure 7, we ensure that the value x'_i corresponding to r'_i is not greater than the value x_i corresponding to r_i. Such a gadget enforces the desired inequality $x_i + x_j \geq x_l$. The gadget enforcing the $\leq$-inequality is obtained in a similar way, but here we need to place the gadget in the left side of $\mathcal{P}$.

Orientation gadget. In the addition gadget, we need to copy three pairwise different line segments from the base line into the gadget. Additionally, we require the segments corresponding to the values of x_i, x_j, x_l to be copied into the gadget in a particular order. To make it possible, we create three right-oriented segments corresponding to each variable and one left-oriented.

To ensure that two base line guard segments specify the same value, we use the orientation gadget. This gadget is a slight modification of the addition gadget, where instead of a guard segment r_l, we place a stationary guard corresponding to a value of 2.5. This allows us to ensure that two line segments from the base line, one left-oriented and one right-oriented, specify the same value. Using orientation gadgets we can enforce consistency between all base line guard segments corresponding to the same variable.

Inversion gadget. In the inversion gadget, we have a right-oriented guard segment r_i and a left-oriented guard segment r_j. We construct an umbra Q_u such that guards at r_i, r_j see the whole critical segment of the umbra if and only if the values x_i, x_j specified by the segments satisfy $x_i \cdot x_j \leq 1$. We also construct a nook Q_n which enforces the inequality $x_i \cdot x_j \geq 1$.

A The art gallery problem is in $\exists\mathbb{R}$

In this section we will prove that the art gallery problem is in the complexity class $\exists\mathbb{R}$. Our proof works also for a more general version of the art gallery problem, where the input polygon can have polygonal holes.

Theorem 4. *The art gallery problem is in the complexity class* $\exists\mathbb{R}$.

Proof. Let $(\mathcal{P}, g)$ be an instance of the art gallery problem where the polygon $\mathcal{P}$ has n vertices, each of which has rational coordinates represented by at most B bits. We show how to construct a quantifier-free formula $\Phi := \Phi(\mathcal{P}, g)$ of the first-order theory of the reals such that Φ is satisfiable if and only if $\mathcal{P}$ has a guard set of cardinality g. The formula Φ has length $O(g^7 n^7 B^2) = O(n^{14} B^2)$ and can be computed in polynomial time. It has been our priority to define the formula Φ so that it is as simple as possible to describe. It might be possible to construct an equivalent but shorter formula.

The description of Φ is similar to the formula Ψ that Micha Sharir described to Efrat and Har-Peled [16]

$$\Psi \;:=\; \left[\exists x_1, y_1, \ldots, x_k, y_k \; \forall p_x, p_y : \text{INSIDE-POLYGON}(p_x, p_y) \Longrightarrow \bigvee_{i=1}^{k} \text{SEES}(x_i, y_i, p_x, p_y)\right].$$

For each $i \in \{1, \ldots, k\}$, the variables x_i, y_i represent the position of guard $g_i := (x_i, y_i)$, and $p := (p_x, p_y)$ represents an arbitrary point. The predicate INSIDE-POLYGON(p_x, p_y) tests if the point p is contained in the polygon $\mathcal{P}$, and SEES(x_i, y_i, p_x, p_y) checks if the guard g_i can see the point p. Thus, the formula is satisfiable if and only if there is a guard set of cardinality g. Note that although the implication "$\Longrightarrow$" is not allowed in the first order theory of the reals, we can always substitute "$A \Longrightarrow B$" by "$\neg A \vee B$".

For the purpose of self-containment, we will briefly repeat the construction of the predicates INSIDE-POLYGON(p_x, p_y) and SEES(x_i, y_i, p_x, p_y). The elementary tool is evaluation of the sign of the determinant $\det(\vec{u}, \vec{v})$ of two vectors $\vec{u}, \vec{v}$. Recall that the sign of the expression $\det(\vec{u}, \vec{v})$ determines whether $\vec{v}$ points to the left of $\vec{u}$ (if $\det(\vec{u}, \vec{v}) > 0$), is parallel to $\vec{u}$ (if $\det(\vec{u}, \vec{v}) = 0$), or points to the right of $\vec{u}$ (if $\det(\vec{u}, \vec{v}) < 0$).

We compute a triangulation $\mathcal{T}$ of the polygon $\mathcal{P}$, e.g., using an algorithm from [13], order the vertices of each triangle of $\mathcal{T}$ in the counter-clockwise order, and orient each edge of the triangle accordingly. A point is contained inside the polygon if and only if it is contained in one of the triangles of $\mathcal{T}$. A point is contained in a triangle if and only if it is on one of the edges or to the left of each edge. Thus the predicate INSIDE-POLYGON(p_x, p_y) has length $O(nB)$.

A guard g_i sees a point p if and only if no two consecutive edges of $\mathcal{P}$ block the visibility. See Figure 8 on why it is not sufficient to check each edge individually. Given a guard g_i, a point p, and two consecutive edges e_1, e_2 of $\mathcal{P}$, it can be checked by evaluating a constant number of determinants whether e_1, e_2 block the visibility between g_i and p. Thus SEES(x_i, y_i, p_x, p_y) has length $O(nB)$ and consequently $\bigvee_{i=1}^{k}$ SEES(x_i, y_i, p_x, p_y) has length $O(knB)$.

Note that the formula Ψ is not a formula in ETR because of the universal quantifier. The main idea to get an equivalent formula with no universal quantifier is to find a polynomial number of points inside $\mathcal{P}$ which, if all seen by the guards, ensure that all of $\mathcal{P}$ is seen. We denote such a set of points as a *witness set*.

Creating a witness set. We are now ready to describe the witness set that replaces the universal quantifier. Let $\mathcal{L} := \{\ell_1, \ldots, \ell_m\}$ be the set of lines containing either an edge of $\mathcal{P}$, or a guard $g \in G$ and a vertex $v \in P$.[†] The well-defined lines in $\mathcal{L}$ partition the plane into *regions*, which are connected components of $\mathbb{R}^2 \setminus \bigcup_{\ell \in \mathcal{L}} \ell$ (see Figure 8 for an example).

[†]Note that if a line is defined as passing through a guard $g \in G$ and a vertex $v \in P$ such that g and v are coincident, the line is not well-defined. Such lines are not considered in the partition into regions described below, but they are included in the set $\mathcal{L}$. Later we will show how to ignore these lines in our formula.

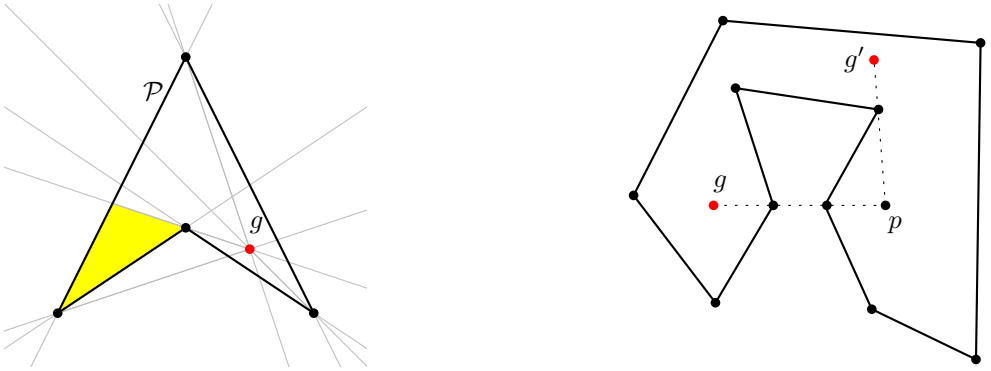

Figure 8: Illustrations for the proof that the art gallery problem is in $\exists\mathbb{R}$. Left: A polygon $\mathcal{P}$ together with a guard g and the resulting line arrangement $\mathcal{L}$. The yellow region is not seen. Right: A line segment gp going through the vertices of $\mathcal{P}$. This illustrates that it is not sufficient to check each edge individually for crossing. A vertex of $\mathcal{P}$ is on the line segment $g'p$, but the visibility between p and g' is not blocked in that case.

Let $\mathcal{A}$ be the set of these regions. It is easy to verify that $\mathcal{A}$ has the following properties:

- Each region in $\mathcal{A}$ is an open convex polygon.
- Each region in $\mathcal{A}$ is either contained in $\mathcal{P}$ or contained in the complement of $\mathcal{P}$.
- The closure of the union of the regions that are contained in $\mathcal{P}$ equals $\mathcal{P}$.
- For each region $R \in \mathcal{A}$, each guard $g \in G$ either sees all points of R or sees no point in R. In particular, if a guard g sees one point in R, it sees all of R and its closure.

Thus it is sufficient to test that for each region $R \in \mathcal{A}$ which is in $\mathcal{P}$, at least one point in R is seen by a guard. For three points a, b, c, define the *centroid* of a, b, c to be the point $C(a, b, c) := (a + b + c)/3$. If a, b, c are three different vertices of the same region $R \in \mathcal{A}$, then the centroid $C(a, b, c)$ must lie in the interior of R. Note that each region has at least three vertices and thus contains at least one such centroid. Let X be the set of all intersection points between two non-parallel well-defined lines in $\mathcal{L}$, i.e., X consists of all vertices of all regions in $\mathcal{A}$. In the formula Φ, we generate all points in X. For any three points a, b, c in X, we also generate the centroid $C(a, b, c)$. If the centroid is in $\mathcal{P}$, we check that it is seen by a guard. Since there are $O(((kn)^2)^3)$ centroids of three points in X and each is tested by a formula of size $O(knB)$, we get a formula of the aforementioned size.

Constructing the formula Φ. Each line ℓ_i is defined by a pair of points $\{(p_i, q_i), (p'_i, q'_i)\}$. Let $\vec{\ell_i} := (p'_i - p_i, q'_i - q_i)$ be a direction vector corresponding to the line. A line ℓ is well-defined if and only if the corresponding vector $\vec{\ell}$ is non-zero.

The lines ℓ_i, ℓ_j are well-defined and non-parallel if and only if $\det(\vec{\ell_i}, \vec{\ell_j}) \neq 0$. If two lines ℓ_i, ℓ_j are well-defined and non-parallel, their intersection point X^{ij} is well-defined and it has coordinates

$$\left(\frac{(p_j q'_j - q_j p'_j)(p'_i - p_i) - (p_i q'_i - q_i p'_i)(p'_j - p_j)}{\det(\vec{\ell_i}, \vec{\ell_j})}, \frac{(p_j q'_j - q_j p'_j)(q'_i - q_i) - (p_i q'_i - q_i p'_i)(q'_j - q_j)}{\det(\vec{\ell_i}, \vec{\ell_j})} \right).$$

For each pair $(i, j) \in \{1, \ldots, m\}^2$, we add the variables x_{ij}, y_{ij} to the formula Φ and we define INTERSECT(i, j) to be the formula

$$\det(\vec{\ell_i}, \vec{\ell_j}) \neq 0 \implies \Big[\det(\vec{\ell_i}, \vec{\ell_j}) \cdot x_{ij} = (p_j q'_j - q_j p'_j)(p'_i - p_i) - (p_i q'_i - q_i p'_i)(p'_j - p_j) \wedge \det(\vec{\ell_i}, \vec{\ell_j}) \cdot y_{ij} = (p_j q'_j - q_j p'_j)(q'_i - q_i) - (p_i q'_i - q_i p'_i)(q'_j - q_j) \Big].$$

It follows that if the formula INTERSECT(i, j) is true then either

- ℓ_i or ℓ_j is not well-defined or they are both well-defined, but parallel, or
- ℓ_i and ℓ_j are well-defined and non-parallel and the variables x_{ij} and y_{ij} are the coordinates of the intersection point X^{ij} of the lines.

Let $\Lambda := \{\lambda_1, \ldots, \lambda_{m^6}\} = \{1, \ldots, m\}^6$ be all the tuples of six elements from the set $\{1, \ldots, m\}$. Each tuple $\lambda := (a, b, c, d, e, f) \in \Lambda$ corresponds to a centroid of the following three points: the intersection point of the lines ℓ_a, ℓ_b, the intersection point of the lines ℓ_c, ℓ_d, and the intersection point of the lines ℓ_e, ℓ_f. For each tuple λ, we proceed as follows. We define the formula CENTROID-DEFINED(λ) to be

$$\det(\vec{\ell_a}, \vec{\ell_b}) \neq 0 \ \wedge \ \det(\vec{\ell_c}, \vec{\ell_d}) \neq 0 \ \wedge \ \det(\vec{\ell_e}, \vec{\ell_f}) \neq 0.$$

We add the variables u_λ, v_λ to the formula Φ, and define the formula CENTROID(λ) as

$$3u_\lambda = x_{ab} + x_{cd} + x_{ef} \ \wedge \ 3v_\lambda = y_{ab} + y_{cd} + y_{ef}.$$

It follows that if the formulas CENTROID-DEFINED(λ) and CENTROID(λ) are both true, then the lines in each of the pairs $(\ell_a, \ell_b), (\ell_c, \ell_d), (\ell_e, \ell_f)$ are well-defined and non-parallel, and the variables u_λ and v_λ are the coordinates of the centroid $C(X^{ab}, X^{cd}, X^{ef})$.

We are now ready to write up our existential formula as

$$\exists x_1, y_1, \ldots, x_k, y_k \ \ \exists x_{11}, y_{11}, x_{12}, y_{12}, \ldots, x_{mm}, y_{mm} \ \ \exists u_{\lambda_1}, v_{\lambda_1}, \ldots, u_{\lambda_{m^6}}, v_{\lambda_{m^6}} : \Phi,$$

where

$$\begin{aligned}
\Phi := & \left[\bigwedge_{(i,j) \in \{1,\ldots,m\}^2} \text{INTERSECT}(i, j) \right] \wedge \\
& \bigwedge_{\lambda \in \Lambda} \Bigg[\text{CENTROID-DEFINED}(\lambda) \implies \\
& \left[\text{CENTROID}(\lambda) \ \wedge \ \left[\text{INSIDE-POLYGON}(u_\lambda, v_\lambda) \implies \bigvee_{i=1}^{k} \text{SEES}(x_i, y_i, u_\lambda, v_\lambda) \right] \right] \Bigg].
\end{aligned}$$

□

B The problem ETR-INV

To show that the art gallery problem is $\exists\mathbb{R}$-hard, we will provide a reduction from the problem ETR-INV, which we introduce below. In this section, we will show that ETR-INV is $\exists\mathbb{R}$-complete.

Definition 5 (ETR-INV)**.** *In the problem ETR-INV, we are given a set of real variables* $\{x_1, \ldots, x_n\}$*, and a set of equations of the form*

$$x = 1, \quad x + y = z, \quad x \cdot y = 1,$$

for $x, y, z \in \{x_1, \ldots, x_n\}$*. The goal is to decide whether the system of equations has a solution when each variable is restricted to the range* $[1/2, 2]$*.*

Thus, if $\Phi(x)$ is an instance of ETR-INV with variables $x := (x_1, \ldots, x_n)$, the space of solutions $S_\Phi := \{x \in [1/2, 2]^n : \Phi(x)\}$ consists of the vectors in $[1/2, 2]^n$ that satisfy all the equations of Φ.

In order to show that ETR-INV is $\exists\mathbb{R}$-complete, we make use of the following problem.

Definition 10. *In the problem* $ETR^{c,+,\cdot}$*, where* $c \in \mathbb{R}$*, we are given a set of real variables* $\{x_1, \ldots, x_n\}$*, and a set of equations of the form*

$$x = c, \quad x + y = z, \quad x \cdot y = z,$$

for $x, y, z \in \{x_1, \ldots, x_n\}$*. The goal is to decide whether the system of equations has a solution.*

A modified version of the problem, where we additionally require that $x_1, \ldots, x_n \in [a, b]$ *for some* $a, b \in \mathbb{R}$*, is denoted by* $ETR^{c,+,\cdot}_{[a,b]}$*.*

We are now ready to prove that ETR-INV is $\exists\mathbb{R}$-complete.

Theorem 6. *The problem ETR-INV is* $\exists\mathbb{R}$*-complete.*

Proof. To show that ETR-INV is $\exists\mathbb{R}$-hard, we will perform a series of polynomial time reductions, starting from ETR and subsequently reducing it to the problems $\mathrm{ETR}^{1,+,\cdot}$, $\mathrm{ETR}^{1/8,+,\cdot}_{[-1/8,1/8]}$, and $\mathrm{ETR}^{1,+,\cdot}_{[1/2,2]}$, and ending with ETR-INV.

To simplify the notation, while considering a problem $\mathrm{ETR}^{c,+,\cdot}$ or $\mathrm{ETR}^{c,+,\cdot}_{[a,b]}$, we might substitute any variable in an equation by the constant c. For instance, $x + c = z$ is a shorthand for the equations $x + y = z$ and $y = c$, where y is an additional variable.

Reduction to $\mathrm{ETR}^{1,+,\cdot}$. We will first argue that $\mathrm{ETR}^{1,+,\cdot}$ is $\exists\mathbb{R}$-hard. This seems to be folklore, but we did not find a formal statement. For the sake of self-containment and rigorousness, we present here a short proof based on the following lemma.

Lemma 11 (Schaefer, Štefankovič [38])**.** *Let* $\Phi(x)$ *be a quantifier-free formula of the first order theory of the reals, where* $x := (x_1, x_2, \ldots, x_n)$ *is a vector of variables. We can construct in polynomial time a polynomial* $F : \mathbb{R}^{n+m} \to \mathbb{R}$ *of degree 4, for some* $m = O(|\Phi|)$*, so that*

$$\{x \in \mathbb{R}^n : \Phi(x)\} = \{x \in \mathbb{R}^n : (\exists y \in \mathbb{R}^m)\ F(x, y) = 0\}.$$

The coefficients of F *have bitlength* $O(|\Phi|)$*.*

Thus it is $\exists\mathbb{R}$-hard to decide if a polynomial has a real root. We reduce this problem to $\mathrm{ETR}^{1,+,\cdot}$. Consider a polynomial equation $Q = 0$. First, we generate all variables corresponding to all the coefficients of Q, by using only the constant 1, addition and multiplication. For example, a variable corresponding to 20 can be obtained as follows: $V_1 = 1, V_{2^1} = 1 + 1, V_{2^2} = V_{2^1} \cdot V_{2^1}, V_{2^4} = V_{2^2} \cdot V_{2^2}, V_{20} = V_{2^4} + V_{2^2}$. We are now left with a polynomial Q' consisting entirely of sums of products of variables, and we keep simplifying Q' as described in the following.

Whenever there is an occurrence of a sum $x + y$ or a product $x \cdot y$ of two variables in Q', we introduce a new variable z. In the first case, we add the equation $x + y = z$ to Φ and substitute the term $x + y$ by z in Q'. In the latter case, we add the equation $x \cdot y = z$ to Φ and substitute $x \cdot y$ by z in Q'. We finish the construction when Q' has been simplified to consist of a single variable, i.e., $Q' = x$, in which case we add the equation $x + V_1 = V_1$ (corresponding to the equation $Q' = 0$) to Φ. When the process finishes, Φ yields an instance of $\mathrm{ETR}^{1,+,\cdot}$, and the solutions to Φ are in one-to-one correspondence with the solutions to the original polynomial equation $Q = 0$.

Reduction to $\mathbf{ETR}^{1/8,+,\cdot}_{[-1/8,1/8]}$. We will now present a reduction from the problem $\mathrm{ETR}^{1,+,\cdot}$ to $\mathrm{ETR}^{1/8,+,\cdot}_{[-1/8,1/8]}$. We use the following result from algebraic geometry, which was stated by Schaefer and Štefankovič [38] in a simplified form. Given an instance $\Phi(x)$ of ETR over the vector of variables $x := (x_1, \ldots, x_n)$, we define the semi-algebraic set S_Φ as the solution space

$$S_\Phi := \{x \in \mathbb{R}^n : \Phi(x)\}.$$

The complexity L of a semi-algebraic set S_Φ is defined as the number of symbols appearing in the formula Φ defining S_Φ (see [27]).

Corollary 12 (Schaefer and Štefankovič [38])**.** *Let B be the set of points in $\mathbb{R}^n$ at distance at most $2^{L^{8n}} = 2^{2^{8n \log L}}$ from the origin. Every non-empty semi-algebraic set S in $\mathbb{R}^n$ of complexity at most $L \geq 4$ contains a point in B.*

Let Φ be an instance of $\mathrm{ETR}^{1,+,\cdot}$ with n variables $x_1, \ldots, x_n$. We construct an instance Φ' of $\mathrm{ETR}^{1/8,+,\cdot}_{[-1/8,1/8]}$ such that Φ has a solution if and only if Φ' has a solution. Let us fix $k := (8n \cdot \log L + 3)$ and $\varepsilon := 2^{-2^k}$. In Φ', we first define a variable V_ε satisfying $V_\varepsilon = \varepsilon$, using $\Theta(k)$ new variables $V_{1/2^{2^2}}, V_{1/2^{2^3}}, \ldots, V_{1/2^{2^k}}$ and equations

$$\begin{aligned} V_{1/2^{2^2}} + V_{1/2^{2^2}} &= 1/8, \\ V_{1/2^{2^2}} \cdot V_{1/2^{2^2}} &= V_{1/2^{2^3}}, \\ V_{1/2^{2^3}} \cdot V_{1/2^{2^3}} &= V_{1/2^{2^4}}, \\ &\vdots \\ V_{1/2^{2^{k-1}}} \cdot V_{1/2^{2^{k-1}}} &= V_\varepsilon. \end{aligned}$$

In Φ', we use the variables $V_{\varepsilon x_1}, \ldots, V_{\varepsilon x_n}$ instead of $x_1, \ldots, x_n$. An equation of Φ of the form $x = 1$ is transformed to the equation $V_{\varepsilon x} = V_\varepsilon$ in Φ'. An equation of Φ of the form $x + y = z$ is transformed to the equation $V_{\varepsilon x} + V_{\varepsilon y} = V_{\varepsilon z}$ of Φ'. An equation of Φ of the form $x \cdot y = z$ is transformed to the following equations of Φ', where $V_{\varepsilon^2 z}$ is a new variable satisfying

$$\begin{aligned} V_{\varepsilon x} \cdot V_{\varepsilon y} &= V_{\varepsilon^2 z}, \\ V_\varepsilon \cdot V_{\varepsilon z} &= V_{\varepsilon^2 z}. \end{aligned}$$

Assume that Φ is true. Then there exists an assignment of values to the variables $x_1, \ldots, x_n$ of Φ that satisfies all the equations and where each variable x_i satisfies $|x_i| \in \left[0, 2^{2^{8n \log L}}\right]$. Then the assignment $V_{\varepsilon x_i} = \varepsilon x_i$ and (when $V_{\varepsilon^2 x_i}$ appears in Φ') $V_{\varepsilon^2 x_i} = \varepsilon^2 x_i$ yields a solution to Φ' with all variables in the range $[-1/8, 1/8]$. On the other hand, if there is a solution to Φ', an analogous argument yields a corresponding solution to Φ. We have given a reduction from $\mathrm{ETR}^{1,+,\cdot}$ to $\mathrm{ETR}^{1/8,+,\cdot}_{[-1/8,1/8]}$. The length of the formula increases by at most a polylogarithmic factor.

Reduction to $\mathbf{ETR}^{1,+,\cdot}_{[1/2,2]}$. We will now show a reduction from $\mathrm{ETR}^{1/8,+,\cdot}_{[-1/8,1/8]}$ to $\mathrm{ETR}^{1,+,\cdot}_{[1/2,2]}$. The reduction is similar as in [42]. We substitute each variable $x_i \in [-1/8, 1/8]$ by $V_{x_i+7/8}$

which will be assumed to have a value of $x_i + 7/8$. Instead of an equation $x = 1/8$ we now have $V_{x+7/8} = 1$. Using addition and the variable equal 1, we can easily get the variables $V_{1/2}, V_{3/2}, V_{3/4}, V_{7/4}, V_{7/8}$ with corresponding values of $1/2, 3/2, 3/4, 7/4$, and $7/8$. Instead of each equation $x + y = z$ we now have equations:

$$\begin{aligned} V_{x+7/8} + V_{y+7/8} &= V_{(z+7/8)+7/8}, \\ V_{z+7/8} + V_{7/8} &= V_{(z+7/8)+7/8}. \end{aligned}$$

As the original variables x, y, z have values in the interval $[-1/8, 1/8]$, the added variables $V_{(z+7/8)+7/8}$ have a value in $[13/8, 15/8]$.

Instead of each equation $x \cdot y = z$ we have the following set of equations

$$\begin{aligned} V_{x+7/8} + V_{y+7/8} &= V_{x+y+14/8}, \quad (V_{x+y+14/8} \in [12/8, 2]) \\ V_{x+y+7/8} + V_{7/8} &= V_{x+y+14/8}, \quad (V_{x+y+7/8} \in [5/8, 9/8]) \\ V_{x+7/8} + V_1 &= V_{x+15/8}, \quad (V_{x+15/8} \in [14/8, 2]) \\ V_{x+1} + V_{7/8} &= V_{x+15/8}, \quad (V_{x+1} \in [7/8, 9/8]) \\ V_{y+7/8} + V_1 &= V_{y+15/8}, \quad (V_{y+15/8} \in [14/8, 2]) \\ V_{y+1} + V_{7/8} &= V_{y+15/8}, \quad (V_{y+1} \in [7/8, 9/8]) \\ V_{x+1} \cdot V_{y+1} &= V_{xy+x+y+1}, \quad (V_{xy+x+y+1} \in [49/64, 81/64]) \\ V_{xy+x+y+1} + V_{1/2} &= V_{xy+x+y+3/2}, \quad (V_{xy+x+y+3/2} \in [81/64, 113/64]) \\ V_{xy+5/8} + V_{x+y+7/8} &= V_{xy+x+y+3/2}, \quad (V_{xy+5/8} \in [39/64, 41/64]) \\ V_{xy+5/8} + 1 &= V_{xy+13/8}, \quad (V_{xy+13/8} \in [103/64, 105/64]) \\ V_{z+7/8} + V_{3/4} &= V_{xy+13/8}. \end{aligned}$$

Each formula Φ of $\mathrm{ETR}^{1/8,+,\cdot}_{[-1/8,1/8]}$ is transformed to a formula Φ' of $\mathrm{ETR}^{1,+,\cdot}_{[1/2,2]}$, as explained above. If there is a solution for Φ, there clearly is a solution for Φ', as all the newly introduced variables have a value within the intervals claimed above. If there is a solution for Φ', there is also a solution for Φ, as the newly introduced variables $V_{2x+7/4}, V_{x+5/8} \in [1/2, 2]$ ensure that $x \in [-1/8, 1/8]$. The increase in the length of the formula is linear.

Note that the only place where we use multiplication is in the formula $V_{x+1} \cdot V_{y+1} = V_{xy+x+y+1}$, where $V_{x+1}, V_{y+1} \in [7/8, 9/8]$. We will use this fact in the next step of the reduction.

Reduction to ETR-INV. We will now show that $\mathrm{ETR}^{1/8,+,\cdot}_{[-1/8,1/8]}$ reduces to ETR-INV. In the first step, we reduce a formula Φ of $\mathrm{ETR}^{1/8,+,\cdot}_{[-1/8,1/8]}$ to a formula Φ' of $\mathrm{ETR}^{1,+,\cdot}_{[1/2,2]}$, as described in the step above. We now have to show how to express each equation $x \cdot y = z$ of Φ' using only the equations allowed in ETR-INV. Note that, as explained above, multiplication is used only for variables $x, y \in [7/8, 9/8]$.

We first show how to define a new variable V_{x^2} satisfying $V_{x^2} = x^2$, where $x \in [7/8, 9/8]$.

$$\begin{aligned} x + V_{3/4} &= V_{x+3/4}, \quad (V_{x+3/4} \in [13/8, 15/8]) \\ V_{1/(x+3/4)} \cdot V_{x+3/4} &= 1, \quad (V_{1/(x+3/4)} \in [8/15, 8/13]) \\ V_{x-1/4} + 1 &= V_{x+3/4}, \quad (V_{x-1/4} \in [5/8, 7/8]) \\ V_{1/(x-1/4)} \cdot V_{x-1/4} &= 1, \quad (V_{1/(x-1/4)} \in [8/7, 8/5]) \\ V_{1/(x^2+x/2-3/16)} + V_{1/(x+3/4)} &= V_{1/(x-1/4)}, \quad (V_{1/(x^2+x/2-3/16)} \in [64/105, 64/65]) \\ V_{1/(x^2+x/2-3/16)} \cdot V_{x^2+x/2-3/16} &= 1, \quad (V_{x^2+x/2-3/16} \in [65/64, 105/64]) \\ x + V_{7/8} &= V_{x+7/8}, \quad (V_{x+7/8} \in [14/8, 2]) \end{aligned}$$

$$
\begin{aligned}
V_{x+1/8} + V_{3/4} &= V_{x+7/8}, \quad (V_{x+1/8} \in [1, 10/8]) \\
V_{x/2+1/16} + V_{x/2+1/16} &= V_{x+1/8}, \quad (V_{x/2+1/16} \in [1/2, 10/16]) \\
V_{x^2-1/4} + V_{x/2+1/16} &= V_{x^2+x/2-3/16}, \quad (V_{x^2-1/4} \in [33/64, 65/64]) \\
V_{x^2-1/4} + V_{3/4} &= V_{x^2+1/2}, \quad (V_{x^2+1/2} \in [81/64, 113/64]) \\
V_{x^2} + V_{1/2} &= V_{x^2+1/2}.
\end{aligned}
$$

Note that the constructed variables are in the range $[\frac{1}{2}, 2]$. In the following, as shorthand for the construction given above, we allow to use equations of the form $x^2 = y$, for a variable x with a value in $[7/8, 9/8]$. We now describe how to express an equation $x \cdot y = z$, where $x, y \in [7/8, 9/8]$.

$$
\begin{aligned}
x + V_{7/8} &= V_{x+7/8}, \quad (V_{x+7/8} \in [14/8, 2]) \\
V_{(x+7/8)/2} + V_{(x+7/8)/2} &= V_{x+7/8}, \quad (V_{(x+7/8)/2} \in [14/16, 1]) \\
y + V_{7/8} &= V_{y+7/8}, \quad (V_{y+7/8} \in [14/8, 2]) \\
V_{(y+7/8)/2} + V_{(y+7/8)/2} &= V_{y+7/8}, \quad (V_{(y+7/8)/2} \in [14/16, 1]) \\
V_{(x+7/8)/2} + V_{(y+7/8)/2} &= V_{(x+y)/2+7/8}, \quad (V_{(x+y)/2+7/8} \in [14/8, 2]) \\
V_{(x+y)/2} + V_{7/8} &= V_{(x+y)/2+7/8}, \quad (V_{(x+y)/2} \in [7/8, 9/8]) \\
V_{(x+y)/2}^2 &= V_{((x+y)/2)^2}, \quad (V_{((x+y)/2)^2} \in [49/64, 81/64]) \\
V_{((x+y)/2)^2} + V_{1/2} &= V_{((x+y)/2)^2+1/2}, \quad (V_{((x+y)/2)^2+1/2} \in [81/64, 113/64]) \\
x^2 &= V_{x^2}, \quad (V_{x^2} \in [49/64, 81/64]) \\
y^2 &= V_{y^2}, \quad (V_{y^2} \in [49/64, 81/64]) \\
V_{x^2} + V_{1/2} &= V_{x^2+1/2}, \quad (V_{x^2+1/2} \in [81/64, 113/64]) \\
V_{x^2/2+1/4} + V_{x^2/2+1/4} &= V_{x^2+1/2}, \quad (V_{x^2/2+1/4} \in [81/128, 113/128]) \\
V_{y^2} + V_{1/2} &= V_{y^2+1/2}, \quad (V_{y^2+1/2} \in [81/64, 113/64]) \\
V_{y^2/2+1/4} + V_{y^2/2+1/4} &= V_{y^2+1/2}, \quad (V_{y^2/2+1/4} \in [81/128, 113/128]) \\
V_{x^2/2+1/4} + V_{y^2/2+1/4} &= V_{(x^2+y^2)/2+1/2}, \quad (V_{(x^2+y^2)/2+1/2} \in [81/64, 113/64]) \\
V_{(x^2+y^2)/2} + V_{1/2} &= V_{(x^2+y^2)/2+1/2}, \quad (V_{(x^2+y^2)/2} \in [49/64, 81/64]) \\
V_{(x^2+y^2)/4+1/4} + V_{(x^2+y^2)/4+1/4} &= V_{(x^2+y^2)/2+1/2}, \quad (V_{(x^2+y^2)/4+1/4} \in [81/128, 113/128]) \\
V_{(x^2+y^2)/4+1/4} + V_{xy/2+1/4} &= V_{((x+y)/2)^2+1/2}, \quad (V_{xy/2+1/4} \in [81/128, 113/128]) \\
V_{xy/2+1/4} + V_{xy/2+1/4} &= V_{xy+1/2}, \quad (V_{xy+1/2} \in [81/64, 113/64]) \\
z + V_{1/2} &= V_{xy+1/2}.
\end{aligned}
$$

The constructed variables are in a range $[1/2, 2]$.

A formula Φ of $\mathrm{ETR}^{1/8,+,\cdot}_{[-1/8,1/8]}$ has been first transformed into a formula Φ' of $\mathrm{ETR}^{1,+,\cdot}_{[1/2,2]}$, and subsequently into a formula Φ'' of ETR-INV. If Φ is satisfiable, then both Φ' and Φ'' are satisfiable. If Φ'' is satisfiable, then both Φ' and Φ are satisfiable. We get that ETR-INV is $\exists\mathbb{R}$-hard.

As the conjunction of the equations of ETR-INV, together with the inequalities describing the allowed range of the variables within ETR-INV, is a quantifier-free formula of the first-order theory of the reals, ETR-INV is in $\exists\mathbb{R}$, which yields that ETR-INV is $\exists\mathbb{R}$-complete. □

Lemma 11 and Corollary 12 together with the reductions explained in this section imply the following lemma.

Lemma 13. *Let Φ be an instance of ETR with variables $x_1, \ldots, x_n$. Then there exists an instance Ψ of ETR-INV with variables $y_1, \ldots, y_m$, $m \geq n$, and constants $c_1, \ldots, c_n, d_1, \ldots, d_n \in \mathbb{Q}$, such that*

- *there is a solution to* Φ *if and only if there is a solution to* Ψ*, and*
- *for any solution* $(y_1, \ldots, y_m)$ *to* Ψ*, there exists a solution* $(x_1, \ldots, x_n)$ *to* Φ *where* $y_1 = c_1 x_1 + d_1, \ldots, y_n = c_n x_n + d_n$*.*

C Reduction from ETR-INV to the art gallery problem

C.1 Notation

Given two different points p, q, the line containing p and q is denoted as $\overleftrightarrow{pq}$, the ray with the origin at p and passing through q is denoted as $\overrightarrow{pq}$, and the line segment from p to q is denoted as pq. For a point p, we let $x(p)$ and $y(p)$ denote the x- and y-coordinate of p, respectively. Table 1 shows the definitions of some objects and distances frequently used in the description of the construction.

C.2 Overview of the construction

Let Φ be an instance of the problem ETR-INV with n variables $X := \{x_0, \ldots, x_{n-1}\}$ and consisting of k equations. We show that there exists a polygon $\mathcal{P} := \mathcal{P}(\Phi)$ with corners at rational coordinates which can be computed in polynomial time such that Φ has a solution if and only if $\mathcal{P}$ can be guarded by some number $g := g(\Phi)$ of guards. The number g will follow from the construction. A sketch of the polygon $\mathcal{P}$ is shown in Figure 9.

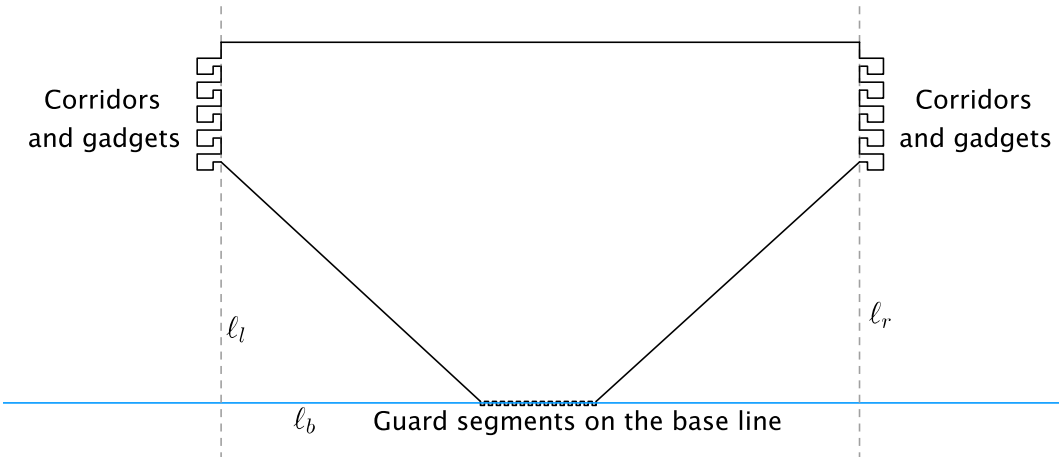

Figure 9: A high-level sketch of the construction of the polygon $\mathcal{P}$.

Each variable $x_i \in X$ is represented by a collection of *guard segments*, which are horizontal line segments contained in the interior of $\mathcal{P}$. Consider one guard segment $s := ab$, where a is to the left of b, and assume that s represents the variable x_i and that there is exactly one guard p placed on s. The guard segment s can be oriented to the right or to the left. The guard p on s specifies the value of the variable x_i as $\frac{1}{2} + \frac{3\|ap\|}{2\|ab\|}$ if s is oriented to the right, and $\frac{1}{2} + \frac{3\|bp\|}{2\|ab\|}$ if s is oriented to the left. Here, the additive term $\frac{1}{2}$ and the factor $\frac{3}{2}$ stem from the fact that all the variables in X are contained in the interval $[\frac{1}{2}, 2]$.

Suppose that there is a solution to Φ. We will show that in that case any minimum guard set G of $\mathcal{P}$ has size $g(\Phi)$ and *specifies* a solution to Φ in the sense that it satisfies the following two properties.

- Each variable $x_i \in X$ is specified *consistently* by G, meaning that there is exactly one guard on each guard segment representing x_i, and all these guards specify the same value of x_i.
- The guard set G is *feasible*, meaning that the values of X thus specified is a solution to Φ.

Moreover, if there is no solution to Φ, each guard set of $\mathcal{P}$ consists of more than $g(\Phi)$ guards.

Name	Description/value
Φ	ETR-INV formula that we reduce from
$X := \{x_0, \ldots, x_{n-1}\}$	set of variables of Φ
$\mathcal{P} := \mathcal{P}(\Phi)$	final polygon to be constructed from Φ
$g := g(\Phi)$	number of guards needed to guard $\mathcal{P}$ if and only if Φ has a solution
k	number of equations in Φ
n	number of variables in Φ
k'	$3n$ plus number of equations $x \cdot y = 1$, $x + y = z$ in Φ
N	$\max\{4n, k'\}$
C	200000
ℓ_b (base line)	line that contains $4n$ guard segments at the bottom of $\mathcal{P}$
$s_i = a_i b_i$	guard segment on ℓ_b, $i \in \{0, \ldots, 4n-1\}$, $\|a_i b_i\| = 3/2$
$\mathcal{P}_M$ (main part of $\mathcal{P}$)	middle part of $\mathcal{P}$, without corridors and gadgets
ℓ_r, ℓ_l	vertical lines bounding $\mathcal{P}_M$
corridor	connection between the main part of $\mathcal{P}$ and a gadget
c_0d_0, c_1d_1	corridor entrances, $d_0 := c_0 + (0, \frac{3}{CN^2})$ and $d_1 := c_1 + (0, \frac{1.5}{CN^2})$
r_i, r_j, r_l, r'_i	guard segments within gadgets, of length $\frac{1.5}{CN^2}$
a'_σ, b'_σ, $\sigma \in \{i, j, l\}$	left and right endpoint of r_σ
ℓ_c	vertical line through $\frac{c_1+c_2}{2}$
o,o'	intersections of rays $\overrightarrow{a_0c_0}$ and $\overrightarrow{b'_lc_1}$ with the line ℓ_c
δ, ρ , ε	$\delta := \frac{13.5}{CN^2}$, $\rho := \frac{\delta}{9} = \frac{1.5}{CN^2}$, $\varepsilon := \frac{\rho}{12} = \frac{1}{8CN^2}$
V	$\frac{c_0+c_1}{2} + (0, 1) + [-38N\rho, +38N\rho] \times [-38N\rho, +38N\rho]$
slab $S(q, v, r)$	region of all points with distance at most r to the line through q with direction v
center of slab	line in the middle of a slab
L-slabs, R-slabs	uncertainty regions for visibility rays, see page 31

Table 1: Parameters, variables, and certain distances that are frequently used are summarized in this table for easy access. Some descriptions are much simplified.

The polygon $\mathcal{P}$ is constructed in the following way. The bottom part of the polygon consists of a collection of *pockets*, containing in total $4n$ collinear and equidistant guard segments. We denote the horizontal line containing these guard segments as the *base line* or ℓ_b. In order from left to right, we denote the guard segments as $s_0, \ldots, s_{4n-1}$. The segments $s_0, \ldots, s_{n-1}$ are right-oriented segments representing the variables $x_0, \ldots, x_{n-1}$, as are the segments $s_n, \ldots, s_{2n-1}$, and $s_{2n}, \ldots, s_{3n-1}$. The segments $s_{3n}, \ldots, s_{4n-1}$ are left-oriented ant they also represent the variables $x_0, \ldots, x_{n-1}$. At the left and at the right side of $\mathcal{P}$, there are some *corridors* attached, each of which leads into a *gadget*. The *entrances* to the corridors at the right side of $\mathcal{P}$ are line segments contained in a vertical line ℓ_r. Likewise, the entrances to the corridors at the left side of $\mathcal{P}$ are contained in a vertical line ℓ_l. The gadgets also contain guard segments, and they are used to impose dependencies between the guards in order to ensure that if there is a solution to Φ, then any minimum guard set of $\mathcal{P}$ consists of $g(\Phi)$ guards and specifies a solution to Φ in the sense defined above. The corridors are used to copy the positions of guards on guard segments on the base line to guards on guard segments inside the gadgets. Each gadget corresponds to a constraint of one of the types $x + y \geq z$, $x + y \leq z$, $x \cdot y = 1$, $x + y \geq 5/2$, and $x + y \leq 5/2$. The first three types of constraints are used to encode the dependencies between the variables in X as specified by Φ, whereas the latter two constraints are used to encode the dependencies between the right-oriented and left-oriented guard segments representing a single variable in X.

C.3 Creating a stationary guard position

We denote some points of $\mathcal{P}$ as *stationary guard positions*. A guard placed at a stationary guard position is called a *stationary guard*. We will often define a stationary guard position as the unique point $p \in \mathcal{P}$ such that a guard placed at p can see some two vertices q_1, q_2 of the polygon $\mathcal{P}$.

We will later prove that for any guard set of size of at most $g(\Phi)$, there is a guard placed at each stationary guard position. For that, we will need the lemma stated below. For an example of the application of the lemma, see Figure 12. The stationary guard position g_2 is the only point from which a guard can see both vertices q_1 and q_2. Applying Lemma 14 with $p := g_2$, $W := \{q_1, q_2\}$, $A := P$ and $M := \{t_1, t_2\}$, we get that there must be a guard placed at g_2 in any guard set of size 3. The purpose of the area A is so that we can restrict our arguments to a small area of the polygon.

Lemma 14. *Let P be a polygon, $A \subseteq P$, and M a set of points in A such that no point in M can be seen from a point in $P \setminus A$, and no two points in M can be seen from the same point in P. Suppose that there is a point $p \in A$ and a set of points $W \subset A$ such that*

1. *no point in W can be seen from a point in $P \setminus A$,*
2. *the only point in P that sees all points in W is p, and*
3. *no point in P can see a point in M and a point in W simultaneously.*

Then any guard set of P has at least $|M| + 1$ guards placed within A, and if a guard set with $|M| + 1$ guards placed within A exists, one of its guards is placed at p.

Proof. Let q be a point in W. Since no two points in the set $\{q\} \cup M$ can be seen from the same point in P, and no point from $P \setminus A$ can see a point in $\{q\} \cup M$, at least $|M| + 1$ guards are needed within A. Suppose that a guard set with exactly $|M| + 1$ guards placed within A exists. There must be $|M|$ guards in A such that each of them can see one point in M and no point in W. The last guard in A has to be at the point p in order to see all points in W. □

In the polygon $\mathcal{P}$ we often use stationary guards for the purpose of seeing some region on one side of a line segment ℓ, but no points on the other side of ℓ. Other guards have the responsibility to see the remaining area. See Figure 10 (left) for an explanation of how a stationary guard position can be constructed.

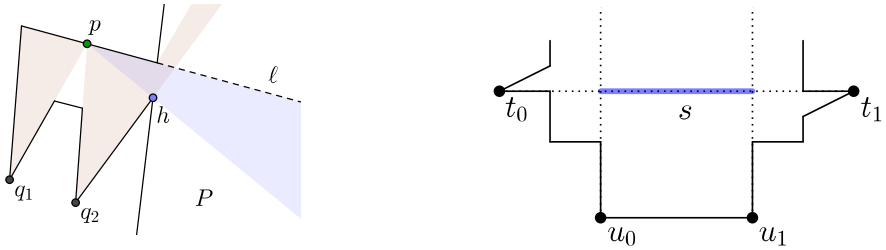

Figure 10: Left: The construction of a stationary guard position p that sees an area in P below a line segment ℓ. The brown areas are the regions of points that see q_1 and q_2, and p is the only point that sees both q_1 and q_2. The point p sees the points in the blue wedge, and the angle of the wedge can be adjusted by choosing the point h accordingly. Right: The construction of a guard segment s (the blue segment). In order to see the points t_0, t_1, a guard must be on the horizontal dotted segment. Furthermore, in order to see u_0, u_1, the guard must be between the vertical dotted segments that contain the endpoints of s. Thus, a guard sees t_0, t_1, u_0, u_1 if and only if the guard is at s.

C.4 Creating a guard segment

In the construction of $\mathcal{P}$ we will denote some horizontal line segments of $\mathcal{P}$ as *guard segments*. We will later prove that for any guard set of size of at most $g(\Phi)$, there is exactly one guard placed at each guard segment.

We will always define a guard segment s by providing a collection of four vertices of $\mathcal{P}$ such that a guard within $\mathcal{P}$ can see all these four vertices if and only if it is placed on the line segment s. See Figure 10 (right) for an example of such a construction. To show that there is a guard placed on a guard segment, we will use the following lemma.

Lemma 15. *Let P be a polygon, $A \subseteq P$, and M a set of points in A such that no point in M can be seen from a point in $P \setminus A$, and no two points in M can be seen from the same point in P. Suppose that there is a line segment s in A, and points $t_0, t_1, u_0, u_1 \in A$ such that*

1. *no point in $\{t_0, t_1, u_0, u_1\}$ can be seen from a point in $P \setminus A$,*
2. *a guard in P sees all of the points t_0, t_1, u_0, u_1 if and only if the guard is at s, and*
3. *no point in P can see a point in M and one of the points t_0, t_1, u_0, u_1.*

Then any guard set of P has at least $|M| + 1$ guards placed within A, and if a guard set with $|M| + 1$ guards placed within A exists, one of its guards is placed on the line segment s.

Proof. Similar to the proof of Lemma 14. □

Consider once more the example pictured in Figure 12, where we want to guard the polygon with only three guards. We define two guard segments, a_0b_0 and a_1b_1. The first one is defined by the vertices t_0, t_1, u_0, u_1, and the second one by the vertices t_2, t_3, u_2, u_3. Applying Lemma 15 we get that there must be a guard placed at a_0b_0 (we set $A := P$ and $M := \{t_2, q_2\}$) and at a_1b_1 (we set $A := P$ and $M := \{t_1, q_2\}$) in any guard set of size 3.

As already explained in Section C.2, guards placed at the guard segments will be used to encode the values of the variables of Φ.

C.5 Imposing inequalities by nooks and umbras

In this section we introduce nooks and umbras, which are our basic tools used to impose dependency between guards placed on two different guard segments. For the following definitions, see Figure 11.

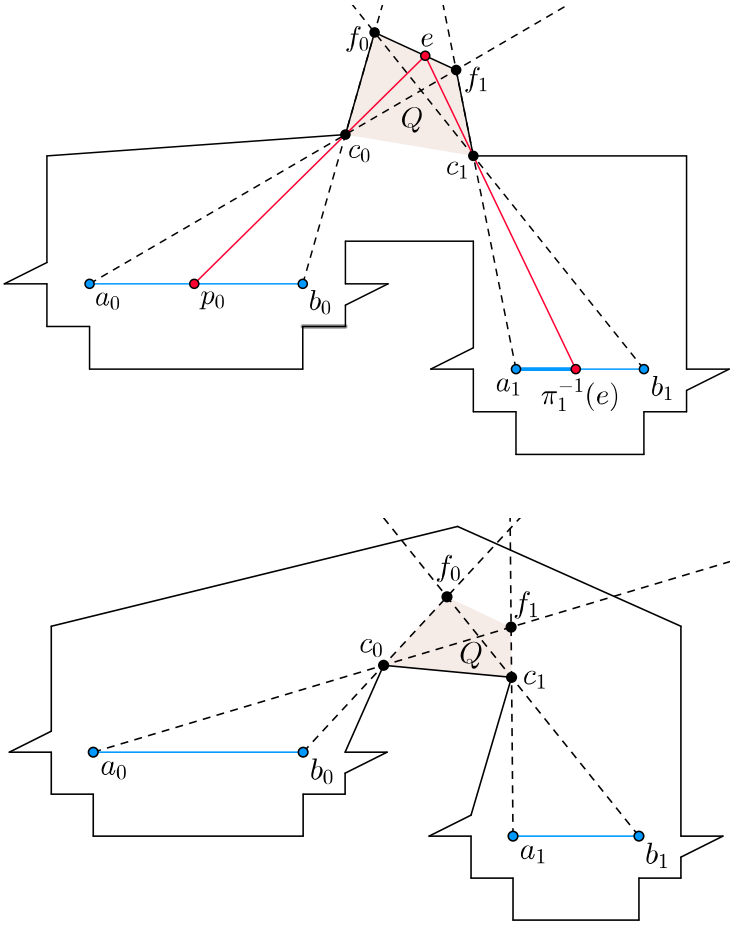

Figure 11: The brown area Q representing a nook (top), and an umbra (bottom). In the left figure, note that if a guard p_1 placed at the segment a_1b_1 has to see the whole line segment f_0f_1 together with p_0, then p_1 must be on or to the left of the point $\pi_1^{-1}(e)$, where $e := \pi_0(p_0)$.

Definition 7 (nook and umbra)**.** *Let $\mathcal{P}$ be a polygon with guard segments $r_0 := a_0b_0$ and $r_1 := a_1b_1$, where r_0 is to the left of r_1. Let c_0, c_1 be two vertices of $\mathcal{P}$, such that c_0 is to the left of c_1. Suppose that the rays $\overrightarrow{b_0c_0}$ and $\overrightarrow{b_1c_1}$ intersect at a point f_0, the lines $\overrightarrow{a_0c_0}$ and $\overrightarrow{a_1c_1}$ intersect at a point f_1, and that $Q := c_0c_1f_1f_0$ is a convex quadrilateral contained in $\mathcal{P}$. For each $i \in \{0,1\}$ define the function $\pi_i \colon r_i \longrightarrow f_0f_1$ such that $\pi_i(p)$ is the intersection of the ray $\overrightarrow{pc_i}$ with the line segment f_0f_1, and suppose that π_i is bijective.*

We say that Q is a nook *of the pair of guard segments r_0, r_1 if for each $i \in \{0,1\}$ and every $p \in r_i$, a guard at p can see all of the segment $\pi_i(p)f_{1-i}$ but nothing else of f_0f_1. We say that Q is an* umbra *of the segments r_0, r_1 if for each $i \in \{0,1\}$ and every $p \in r_i$, a guard at p can see all of the segment $\pi_i(p)f_i$ but nothing else of f_0f_1. The functions π_0, π_1 are called* projections *of the nook or the umbra.*

We will construct nooks and umbras for pairs of guard segments where we want to enforce dependency between the values of the corresponding variables. When making use of an umbra, we will also create a stationary guard position from which a guard sees the whole quadrilateral Q, but nothing on the other side of the line segment f_0f_1. In this way we can enforce the guards on r_0 and r_1 to see all of f_0f_1 together. For the case of a nook, the segment f_0f_1 will always be on the polygon boundary, and then there will be no stationary guard needed. See Figure 12 for an example of a construction of both a nook and an umbra for a pair of guard segments.

Definition 8 (critical segment and shadow corners)**.** *Consider a nook or an umbra $Q :=$*

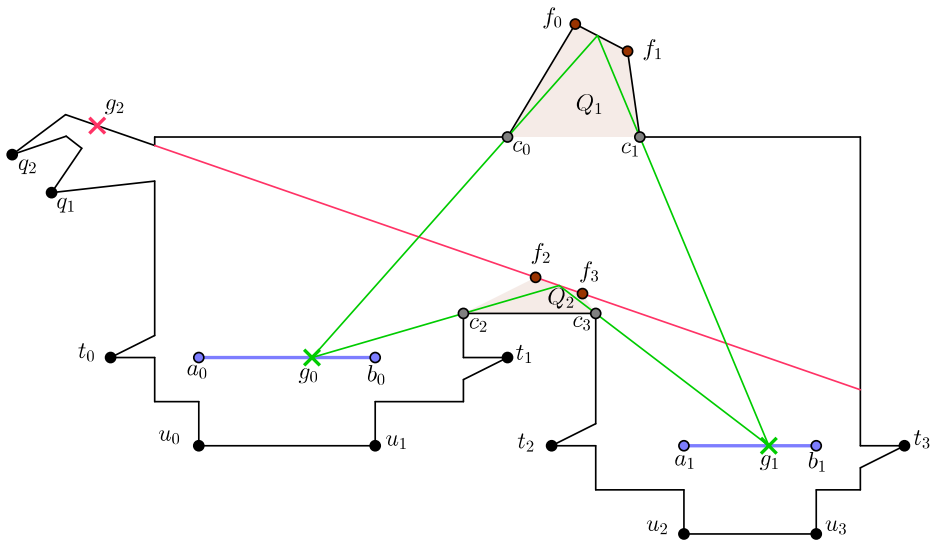

Figure 12: Q_1 is a copy-nook of the segments $r_0 := a_0b_0$ and $r_1 := a_1b_1$ with a critical segment f_0f_1, and Q_2 is a copy-umbra for the same pair with a critical segment f_2f_3. Lemmas 14 and 15 imply that this polygon cannot be guarded by fewer than 3 guards, and any guard set with 3 guards must contain a guard g_0 on r_0, a guard g_1 on r_1, and a stationary guard at the point g_2. The guards g_0 and g_1 must specify the same value on r_0 and r_1, respectively.

$c_0c_1f_1f_0$ of a pair of guard segments r_0, r_1. The line segment f_0f_1 is called the critical segment *of Q, and the vertices c_0, c_1 are called the* shadow corners *of Q.*

Consider a nook or an umbra of a pair of guard segments r_0, r_1. Let p_0, p_1 be the guards placed on the guard segments r_0 and r_1, respectively, and assume that p_0 and p_1 together see all of the critical segment f_0f_1. Let $e := \pi_0(p_0)$. The condition that p_0, p_1 together see all of f_0f_1 enforces dependency between the position of the guard p_1 and the point $\pi_1^{-1}(e)$. If Q is a nook, p_1 must be in the closed wedge W between the rays $\overrightarrow{ec_0}$ and $\overrightarrow{ec_1}$. If Q is an umbra, p_1 must be in $\overline{\mathcal{P} \setminus W}$ (the closure of the complement of W), i.e., either on or to the right of $\pi_1^{-1}(e)$. This observation will allow us to impose an inequality on the x-coordinates of p_0, p_1, and thus on the variables corresponding to the guard segments r_0, r_1.

C.6 Copying one variable

Definition 9. *Let Q be a nook or an umbra of a pair of guard segments $r_0 := a_0b_0$ and $r_1 := a_1b_1$ with the same orientation, such that the shadow corners c_0 and c_1 have the same y-coordinate. We then call Q a* copy-nook *or a* copy-umbra*, respectively.*

We can show the following result.

Lemma 16. *Let Q be a copy-nook or a copy-umbra for a pair of guard segments $r_0 := a_0b_0$ and $r_1 := a_1b_1$. Then for every point $e \in f_0f_1$ we have $\frac{\|a_0\pi_0^{-1}(e)\|}{\|a_0b_0\|} = \frac{\|a_1\pi_1^{-1}(e)\|}{\|a_1b_1\|}$, i.e., the points $\pi_0^{-1}(e)$ and $\pi_1^{-1}(e)$ on the corresponding guard segments r_0 and r_1 represent the same value.*

Proof. See Figure 13. Let $\ell := \overleftrightarrow{c_0c_1}$ be the horizontal line containing the line segment c_0c_1, and ℓ' a horizontal line passing trough e. Let f_2 be an intersection point of the line $\overleftrightarrow{f_0f_1}$ with the line ℓ. Let a_0', a_1', b_0', b_1' be the intersection points of the rays $\overrightarrow{a_0c_0}, \overrightarrow{a_1c_1}, \overrightarrow{b_0c_0}, \overrightarrow{b_1c_1}$, respectively, with the line ℓ'.

We obtain

$$\frac{\|a_0\pi_0^{-1}(e)\|}{\|\pi_0^{-1}(e)b_0\|} \cdot \frac{\|\pi_1^{-1}(e)b_1\|}{\|a_1\pi_1^{-1}(e)\|} = \frac{\|ea_0'\|}{\|b_0'e\|} \cdot \frac{\|b_1'e\|}{\|ea_1'\|} = \frac{\|ea_0'\|}{\|ea_1'\|} \cdot \frac{\|b_1'e\|}{\|b_0'e\|} = \frac{\|c_0f_2\|}{\|c_1f_2\|} \cdot \frac{\|c_1f_2\|}{\|c_0f_2\|} = 1.$$

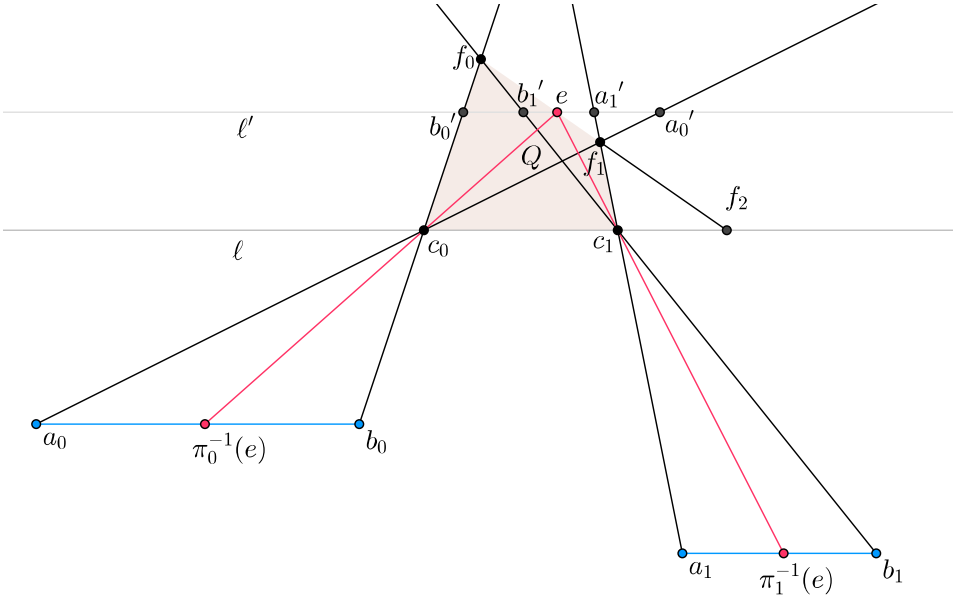

Figure 13: A copy-nook or a copy-umbra Q for a pair of guard segments $r_0 := a_0b_0$ and $r_1 := a_1b_1$. The points $\pi_0^{-1}(e)$ and $\pi_1^{-1}(e)$ represent the same value.

The first equality holds as the following pairs of triangles are similar: $a_0\pi_0^{-1}(e)c_0$ and $a_0'ec_0$, $\pi_0^{-1}(e)b_0c_0$ and $eb_0'c_0$, $a_1\pi_1^{-1}(e)c_1$ and $a_1'ec_1$, $\pi_1^{-1}(e)b_1c_1$ and $eb_1'c_1$. The third equality holds as the following pairs of triangles are similar: $ea_0'f_1$ and $f_2c_0f_1$, $ea_1'f_1$ and $f_2c_1f_1$, $b_1'ef_0$ and $c_1f_2f_0$, and $b_0'ef_0$ and $c_0f_2f_0$. □

The following lemma is a direct consequence of Lemma 16. See Figure 12 for an example of how a construction as described in the lemma can be made.

Lemma 17. *Let r_0, r_1 be a pair of guard segments oriented in the same way for which there is both a copy-nook and a copy-umbra. Suppose that there is exactly one guard p_0 placed at r_0 and one guard p_1 placed at r_1, and that the guards p_0 and p_1 see together both critical segments. Then the guards p_0 and p_1 specify the same value.*

Definition 18. *Let r_0, r_1 be a pair of guard segments for which there is both a copy-nook and a copy-umbra. We say that r_1 is a* copy *of r_0. If there is only a copy-nook* or *a copy-umbra of the pair r_0, r_1, we say that r_1 is a* weak copy *of r_0.*

It will follow from the construction of the polygon $\mathcal{P}$ that if there is a solution to Φ, then for any optimal guard set of $\mathcal{P}$ and any pair of guard segments r_0, r_1 such that r_1 is a copy of r_0, there is exactly one guard on each segment r_0, r_1, and the guards together see the whole critical segment of both the copy-nook and the copy-umbra.

C.7 The overall design of the polygon $\mathcal{P}$

Recall the high-level sketch of the polygon $\mathcal{P}$ in Figure 9. The bottom part of the polygon consists of *pockets* containing $4n$ guard segments $s_0, \ldots, s_{4n-1}$. The guard segments are placed on the base line ℓ_b, each segment having a width of $3/2$ and contained within a pocket of width 13.5. Therefore the horizontal space used for the $4n$ guard segments on the base line ℓ_b is $54n$. The wall of $\mathcal{P}$ forming the $4n$ pockets is denoted the *bottom wall*. The detailed description of the pockets and of the bottom wall is presented in Section C.8.

Let k' be equal to $3n$ plus the number of equations in Φ of the form $x_i + x_j = x_k$ or $x_i \cdot x_j = 1$. At the right side of $\mathcal{P}$ and at the left side of $\mathcal{P}$ there will be at most k' *corridors* attached, each of which leads into a *gadget*. The *entrances* to the corridors are contained in the vertical lines

ℓ_r and ℓ_l. The corridors are described in Section C.9, and they are placed equidistantly, with a vertical distance of 3 between the entrances of two consecutive corridors along the lines ℓ_r and ℓ_l. The gadgets are described in Sections C.10–C.13. The total vertical space occupied by the corridors and the gadgets at each side of $\mathcal{P}$ is at most $3k'$.

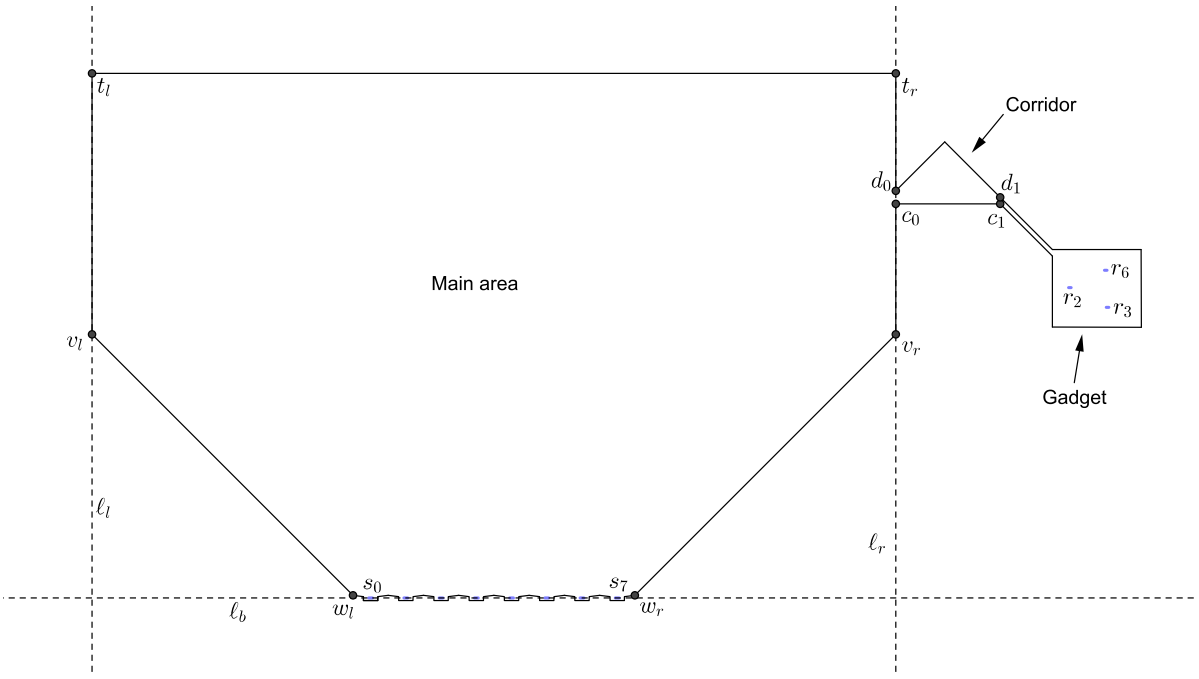

Figure 14: A sketch of the construction of $\mathcal{P}$ with 8 guard segments $s_0, \ldots, s_7$ on the base line, and one gadget at the right side of $\mathcal{P}$. The gadget contains the guard segments r_2, r_3, r_6, which are copies of the segments s_2, s_3, s_6, respectively. The gadget is the part of the polygon to the right of the line segment c_1d_1. The proportions in the drawing are not correct.

Consider the sketch of the polygon $\mathcal{P}$ in Figure 14. Let $N := \max\{4n, k'\}$, and let us define a constant $C := 200000$. Let w_l and w_r denote the left and right endpoint of the bottom wall, respectively. The horizontal distance from w_l to the line ℓ_l is $CN^2 - 54n + 6$, as is the horizontal distance from w_r to ℓ_r. The horizontal distance from the left endpoint a_0 of the leftmost segment to ℓ_r, as well as the horizontal distance from the right endpoint b_{4n-1} of the rightmost segment to ℓ_l, is CN^2. The vertical distance from ℓ_b up to the entrance of the first corridor is CN^2. The boundary of $\mathcal{P}$ contains an edge connecting w_l to the point $v_l := w_l + (-(CN^2 - 54n + 6), CN^2 - 1)$ on ℓ_l, and an edge connecting w_r to the point $v_r := w_r + (CN^2 - 54n + 6, CN^2 - 1)$ on ℓ_r. Let $t_l := v_l + (0, 3k')$ and $t_r := v_r + (0, 3k')$. The *main area* $\mathcal{P}_M$ of $\mathcal{P}$ is the area bounded by the bottom wall of $\mathcal{P}$ (to be defined in Section C.8) and a polygonal curve defined by the points $w_l v_l t_l t_r v_r w_r$. The entrances to the corridors are on the segment $v_l t_l$ in the left side and on the segment $v_r t_r$ in the right side. The set $\mathcal{P} \setminus \mathcal{P}_M$ outside of the main area consists of corridors and gadgets.

The reason why we need the distances from the guard segments on the base line ℓ_b to the gadgets to be so large is that we want all the rays from the guard segments on the base line through the corridor entrances on ℓ_r (ℓ_l) to have nearly the same slopes. That will allow us to describe a general method for copying guard segments from the base line into the gadgets.

Each gadget corresponds to a constraint involving either two or three variables, where each variable corresponds to a guard segment on the base line. Gadgets are connected with the main area $\mathcal{P}_M$ via corridors. A corridor does not contain any guard segments, and its aim is enforcing consistency between (two or three) pairs of guard segments, where one segment from each pair is in $\mathcal{P}_M$ and the other one is in the gadget. Each corridor has two vertical *entrances*, the entrance c_0d_0 of height $\frac{3}{CN^2}$ connecting it with $\mathcal{P}_M$, and the entrance c_1d_1 of height $\frac{1.5}{CN^2}$

connecting it with a gadget. The bottom wall of a corridor is a horizontal line segment c_0c_1 of length 2. The shape of the upper wall is more complicated, and it depends on the indices of the guard segments involved in the corresponding constraint, and on the height of where the corridor is placed with respect to the base line ℓ_b.

A gadget can be thought of as a room which is connected with the main area $\mathcal{P}_M$ of $\mathcal{P}$ via a corridor, i.e., attached to the vertices c_1 and d_1 of the corridor. There are five different kinds of gadgets, each corresponding to a different kind of inequality or equation, and, unlike for the case of corridors, all gadgets of the same type are identical. Each gadget contains one or two guard segments for each variable present in the corresponding formula. All guard segments within a gadget are of length $\frac{1.5}{CN^2}$, and are placed very close to the *middle point* of the gadget, defined as $m := c_1 + (1, -1)$ for gadgets at the right side of $\mathcal{P}$, and $m := c_1 + (-1, -1)$ for gadgets at the left side of $\mathcal{P}$.

C.8 Construction of the bottom wall

In this section we present the construction of the bottom wall of $\mathcal{P}$. We first describe the overall construction, as shown in Figure 15, and later we introduce small features corresponding to each equation of the type $x_i = 1$ in Φ.

The bottom wall forms $4n$ *pockets*, each pocket containing one guard segment on the base line ℓ_b. Each pocket has a width of 13.5. Each guard segment has a length of 3/2, and the distance between two consecutive segments is 12.

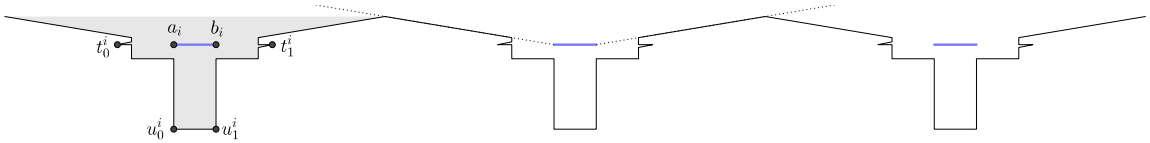

Figure 15: The construction of three consecutive guard segments (blue) on the base line. A pocket corresponding to a single guard segment $s_i := a_ib_i$ is marked in grey.

Let $s_0 := a_0b_0, \ldots, s_{4n-1} := a_{4n-1}b_{4n-1}$ be the guard segments in order from left to right. A pocket for a guard segment $s_i = a_ib_i$ is constructed as shown in Figure 15 (the grey area in the figure). The left endpoint of the pocket is at the point $a_i + (-6, 1)$, and the right endpoint of the pocket is at the point $a_i + (7.5, 1)$. The guard segment is defined by the following points, which are vertices of the pocket: $t_0^i := a_i + (-2, 0)$, $t_1^i := a_i + (3.5, 0)$, $u_0^i := a_i + (0, -5)$, $u_1^i := a_i + (1.5, -5)$. The vertical edges of the pocket are contained in lines $x = x(a_i) - 1.5$, $x = x(a_i)$, $x = x(a_i) + 1.5$, and $x = x(a_i) + 3$. The horizontal edges of the pocket are contained in lines $y = 0, y = -0.5$, and $y = -5$. The remaining edges are constructed so that the points t_0^i, t_1^i can be seen only from within the pocket, and that any point on s_i sees the whole area of the pocket.

Consider an equation of the form $x_i = 1$ in Φ. There are four guard segments representing x_i, i.e., the guard segments s_i, s_{i+n}, s_{i+2n}, and s_{i+3n}, where the first three are right-oriented and the last one is left-oriented. We add two spikes in the construction of the leftmost of these guard segments, i.e., the segment s_i, as shown in Figure 16. The dashed lines in the figure intersect at the point $g_i \in s_i$, where $g_i := a_i + (1/2, 0)$. The spike containing q_1^i enforces the guard to be at the point g_i or to the right of it, while the spike containing q_2^i enforces the guard to be at g_i or to the left of it. Also, the points q_1^i and q_2^i are chosen so that they cannot be seen by any points from within the corridors or gadgets. The guard segment is thus reduced to a stationary guard position g_i corresponding to the value $x_i = 1$.

Note that we only need to add such spikes to the pocket containing s_i, since the construction described in Section C.12 will enforce the guards on all the segments s_i, s_{i+n}, s_{i+2n}, and s_{i+3n} to specify the value of x_i consistently. We have now specified all the details of the main area $\mathcal{P}_M$ of $\mathcal{P}$ (recall the definition of $\mathcal{P}_M$ from Section C.7). The following lemma holds.

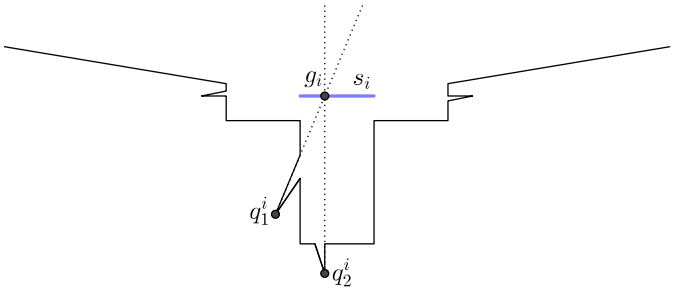

Figure 16: The spikes with vertices at q_1^i and q_2^i enforce the guard from the guard segment s_i to be at the point g_i corresponding to the value of 1.

Lemma 19. *For any ETR-INV formula Φ, we can construct in polynomial time the bottom wall corresponding to Φ such that every vertex has rational coordinates, with the enumerator upper-bounded by ζN and denominator upper-bounded by ζ, where ζ is some universal constant.*

We can prove the following lemma.

Lemma 20. *Any guard set G of the polygon $\mathcal{P}$ satisfies the following properties.*

- *G has at least $4n$ guards placed in $\mathcal{P}_M$.*
- *If G has exactly $4n$ guards placed in $\mathcal{P}_M$, then it has one guard within each of the $4n$ guard segments $s_0, \ldots, s_{4n-1}$.*
- *If G has exactly $4n$ guards placed in the main area of $\mathcal{P}$, then for each variable $x_i \in X$ such that there is an equation $x_i = 1$ in Φ, the guard at the segment s_i is at the position corresponding to the value 1.*

Moreover, any set G' of points such that (i) G' has a point within each of the $4n$ guard segments of $\mathcal{P}_M$, and (ii) for each variable $x_i \in X$ such that there is an equation $x_i = 1$ in Φ there is a guard at s_i at the position corresponding to the value of 1, can see together the whole area of $\mathcal{P}_M$.

Proof. Recall that each point t_1^i, for $i \in \{0, \ldots, 4n-1\}$, can be seen only from within the pocket corresponding to the guard segment s_i. Therefore any guard set requires at least one guard within each of these pockets, i.e., it contains at least $4n$ guards in $\mathcal{P}_M$.

Now, consider any guard segment s_i, for $i \in \{0, \ldots, 4n-1\}$. Let $M := \{t_1^j : j \in \{0, 1, \ldots, 4n-1\}, j \neq i\}$. Note also that none of the points $t_0^i, t_1^i, u_0^i, u_1^i$, and also no point from M can be seen by a guard placed within a corridor or a gadget of $\mathcal{P}$, i.e., outside of the main area of $\mathcal{P}$. By Lemma 15, by taking $t_0^i, t_1^i, u_0^i, u_1^i$ as t_0, t_1, u_0, u_1 and $A := \mathcal{P}_M$, we obtain that if a guard set G has exactly $4n$ guards placed in $\mathcal{P}_M$, then it has a guard within s_i.

For the third property, consider any variable $x_i \in X$ such that there is an equation $x_i = 1$ in Φ. As none of the points q_1^i and q_2^i can be seen from guards within the corridors or gadgets, or from the guards within the other guard segments on the base line, both of them must be seen by the only guard g_i placed on s_i. Then, g_i must be placed at the position corresponding to the value 1.

At last, consider any set G' of points such that (i) G' has a point within each of the $4n$ guard segments of $\mathcal{P}_M$, and (ii) for each variable $x_i \in X$ such that there is an equation $x_i = 1$ in Φ there is a guard at s_i at the position corresponding to the value of 1. We will show that G' can see together the whole area of $\mathcal{P}_M$. From the construction of the pockets, a guard within a pocket containing s_i can see the whole area of the pocket (in particular, the guards at positions corresponding to the value of 1 can see the whole area of the added spikes). The guard at s_i also

sees the whole area of $\mathcal{P}_M$ which is above the pocket, i.e., all points of $\mathcal{P}_M$ with x-coordinates in $[x(a_i) - 6, x(a_i) + 7.5]$. The part of $\mathcal{P}_M$ to the left of the leftmost pocket can be seen by the guard on the leftmost guard segment, and the part of $\mathcal{P}_M$ to the right of the rightmost pocket can be seen by the guard on the rightmost guard segment. □

C.9 Construction of a corridor

In this section, we describe the construction of a corridor. Inside each gadget there are two or three guard segments r_i, r_j, r_l (or r_i, r_j) corresponding to two or three pairwise different guard segments from the base line s_i, s_j, s_l (or s_i, s_j). We require that for each $\sigma \in \{i, j, l\}$ the guard segments s_σ, r_σ have the same orientation. For the corridors attached at the right side of $\mathcal{P}$ we assume $i < j < l$, and for the corridors attached at the left side we assume $i > j > l$. We describe here how to construct a corridor that ensures that that the segments r_i, r_j, r_l are copies of the segments s_i, s_j, s_l, respectively. This construction requires that the guard segments within the gadget satisfy the conditions of some technical lemmas (see Lemma 21 and 26).

Note that this construction can be generalized for copying an arbitrary subset of guard segments, but since we only need to copy two or three segments, we explain the construction in the setting of three segments. The construction for two segments is analogous but simpler. We first describe how to copy into a gadget at the right side of the polygon $\mathcal{P}$ – copying into gadgets at the left side of $\mathcal{P}$ can be done in a symmetric way and is described shortly in Section C.9.4.

As described briefly in Section C.7, the lower wall of the corridor of the gadget is a horizontal edge c_0c_1 of length 2, where c_0 is on the line ℓ_r and c_1 is to the right of c_0. The upper wall of the corridor is more complicated, and it will be described later. It has the left endpoint at $d_0 := c_0 + (0, \frac{3}{CN^2})$, and the right endpoint $d_1 := c_1 + (0, \frac{1.5}{CN^2})$. The vertical line segments c_0d_0 and c_1d_1 are called the *entrances* of the corridor.

C.9.1 Idea of the copying construction

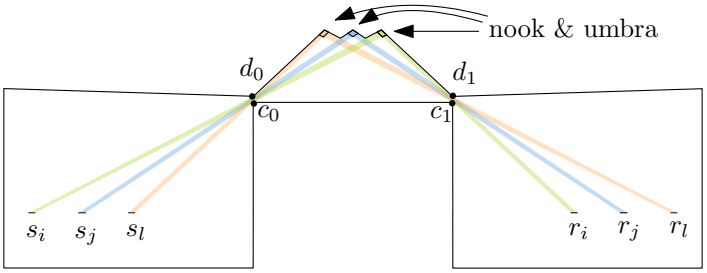

Figure 17: In this figure, we display a simplified corridor construction. The vertices c_0, c_1 serve as shadow corners for three copy-umbras simultaneously for the pairs $(s_i, r_i), (s_j, r_j), (s_l, r_l)$. Each of these pairs also have a small copy-nook in the top of corridor. The entrances c_0d_0 and c_1d_1 to the corridor are sufficiently small so that the critical segments of the nook and umbra of each pair of segments s_σ, r_σ (contained in the small boxes at the top of the figure) are not seen by other guard segments.

To ensure that the segments r_i, r_j, r_l are copies of the segments s_i, s_j, s_l, we need to construct within the corridor copy-nooks and copy-umbras for the pairs of corresponding segments, see Figure 17 for a simplified illustration. The vertices c_0, c_1 of the corridor act as shadow corners in three overlapping copy-umbras for the pairs (s_i, r_i), (s_j, r_j), and (s_l, r_l), respectively. We construct the chain of $\mathcal{P}$ from d_0 to d_1 bounding the corridor from above so that it creates three copy-nooks for the same pairs. To enforce that for any guard set of size $g(\Phi)$, for each $\sigma \in \{i, j, l\}$ the guard segments s_σ and r_σ specify the same value, we have to ensure that no guards on guard segments other than s_σ and r_σ can see the critical segments of the copy-umbra and the copy-nook of the pair s_σ, r_σ. For each $\sigma \in \{i, j, l\}$ we also introduce a stationary

guard position, so that guards placed at these positions together see all the copy-umbras, but nothing on the other sides of the critical segments of the copy-umbras. We also need to ensure that the guards placed at the stationary guard positions cannot see the critical segments of the copy-umbras and the copy-nooks of other pairs. We then obtain (see Lemma 24) that for any guard set with one guard at each guard segment, and with no guards placed outside of the guard segments and stationary guard positions, the segments s_σ, r_σ specify the value of x_σ consistently.

Our construction will ensure that for any $\sigma \in \{i, j, l\}$, only the guard segment s_σ from the base line, and only the guard segment r_σ from within the gadget can see the critical segments of the corresponding copy-nook and copy-umbra. In particular, we will ensure that the vertical edge of $\mathcal{P}$ directly above the entrance c_0d_0 blocks visibility from all guard segments $s_{\sigma'}$ for $\sigma' \in \{0, \ldots, \sigma - 1\}$, whereas the vertical edge of $\mathcal{P}$ directly below c_0d_0 blocks visibility from all guard segments $s_{\sigma'}$ for $\sigma' \in \{\sigma + 1, \ldots, 4n - 1\}$. An analogous property will be ensured for the gadget guard segments.

The main idea to achieve the above property is to make the entrances c_id_i of the corridor sufficiently small. However, we cannot place the point d_0 arbitrarily close to c_0, since both endpoints a_σ and b_σ of the segment s_σ have to see the left endpoint of the critical segment of the copy-umbra, and the right endpoint of the critical segment of the copy-nook for the pair s_σ, r_σ (the points f_0 and f_1, respectively, in the context of Section C.5). By placing the corridor sufficiently far away from the segments on the base line, we obtain that the visibility lines from the guard segment endpoints through the points c_0, d_0 are almost parallel and can be described by a simple pattern. The same holds for the pair of points d_1 and c_1 and the guard segment r_σ. The pattern enables us to construct the corridor with the desired properties.

In the following, we introduce objects that make it possible to describe the upper corridor wall and prove that the construction works as intended.

C.9.2 Introducing slabs

In a small area around the point $\frac{c_0+c_1}{2} + (0, 1)$, every ray from an endpoint of a base line guard segment through one of the points c_0, d_0 intersects every ray from an endpoint of a gadget guard segment through one of the points c_1, d_1. These rays intersect at angles close to $\pi/2$, and they form an arrangement consisting of quadrilaterals, creating a nearly-regular pattern. However, the arrangement of rays is not completely regular. We therefore introduce a collection of thin slabs, where each slab contains one of the rays in a small neighbourhood around $\frac{c_0+c_1}{2} + (0, 1)$, and such that the slabs form an orthogonal grid with axis $(1, 1)$ and $(-1, 1)$. Thus, the slabs are introduced in order to handle the "uncertainty" and irregularity of the rays.

Given a point q and a vector v, the *slab* $S(q, v, r)$ consists of all points at a distance of at most r from the line through q parallel to v. The *center* of the slab $S(q, v, r)$ is the line through q parallel to v.

Let $r_i := a'_ib'_i, r_j := a'_jb'_j, r_l := a'_lb'_l$. Let ℓ_c be a vertical line passing through the middle of the segment c_0c_1. Recall that here we describe the construction of a corridor to be attached at the right side of $\mathcal{P}$. Let o be the intersection point of the ray $\overrightarrow{a_0c_0}$ with ℓ_c, and o' the intersection point of the ray $\overrightarrow{b'_lc_1}$ with ℓ_c. All the points $o, o', \frac{c_0+c_1}{2} + (0, 1)$ lie on the vertical line ℓ_c.

Let us define vectors $\alpha := (1, 1)$, $\beta := (-1, 1)$, and introduce a grid of slabs parallel to α and β. Let us fix $\delta := \frac{13.5}{CN^2}$, $\rho := \frac{\delta}{9} = \frac{1.5}{CN^2}$, and $\varepsilon := \frac{\rho}{12} = \frac{1}{8CN^2}$. For each $\sigma \in \{0, \ldots, 4n - 1\}$ and $\gamma \in \{0, 1, 2, 3\}$ we define a slab

$$L^\gamma_\sigma := S(o + (0, \sigma\delta + \gamma\rho), \alpha, \varepsilon),$$

which we denote as an *L-slab*. Let $\tau(i) := 2$, $\tau(j) := 1$, and $\tau(l) := 0$. For each $\sigma \in \{i, j, l\}$ and $\gamma \in \{0, 1, 2, 3\}$ we define a slab

$$R^\gamma_\sigma := S(o' + (0, \tau(\sigma)\delta + \gamma\rho), \beta, \varepsilon),$$

which we denote as an *R-slab*.

In the case of gadgets with just two guard segments r_i, r_j, we define the point o' as the intersection point of the ray $\overrightarrow{b'_j c_1}$ with ℓ_c, and we define $\tau(i) := 1$ and $\tau(j) := 0$. Then, the R-slabs R^γ_σ are defined as above for $\sigma \in \{i, j\}$.

See Figure 18 for an illustration of the area where the L-slabs intersect the R-slabs.

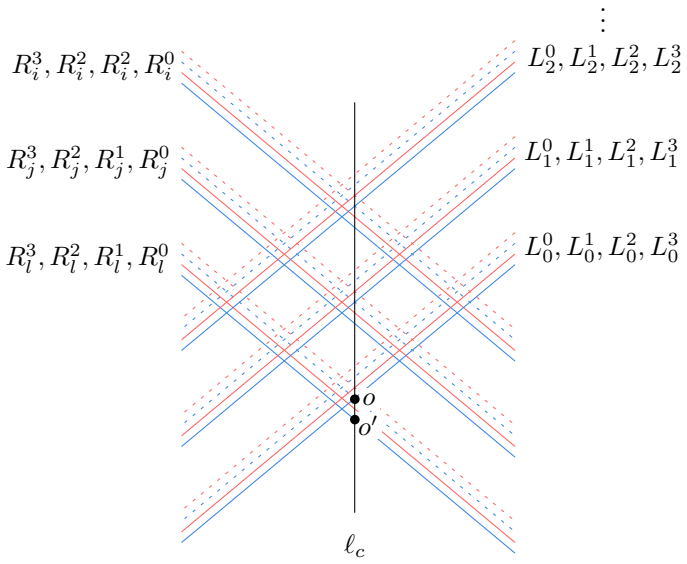

Figure 18: The L-slabs have slope 1 and the R-slabs have slope -1. For each guard segment we get 4 equidistant slabs. The width of each slab is 2ε. The distance between two slabs from the same group is ρ and the distance between two groups is δ. All intersections are contained in the region denoted by V.

Let V be the square $\frac{c_0+c_1}{2} + (0,1) + [-38N\rho, 38N\rho] \times [-38N\rho, 38N\rho]$. We now prove that the area where the L-slabs and the R-slabs intersect is contained in V. Let us denote by $\mathcal{R}$ all the rays with endpoint at one of the guard segments in the main area, going through c_0 and d_0 and all the rays from the endpoints of gadget guard segments through the points c_1, d_1. We also ensure that all rays in $\mathcal{R}$ are inside a predefined slab within the area V.

In sections specific to the particular gadgets, we will prove the following lemma. We state the lemma for gadgets with three guard segments r_i, r_j, r_l, but it has a natural analogy for gadgets with just two guard segments r_i, r_j.

Lemma 21. *For any gadget to be attached at the right side of the polygon $\mathcal{P}$ and containing the guard segments $r_i = a'_i b'_i, r_j = a'_j b'_j, r_l = a'_l b'_l$ the following holds, where c_1 is the bottom-right endpoint of the corridor corresponding to the gadget.*

1. *The intersection of any R-slab with the line ℓ_c is contained in V.*
2. *For each $\sigma \in \{i, j, l\}$, it holds that $\overrightarrow{b'_\sigma c_1} \cap V \subset R^0_\sigma, \overrightarrow{a'_\sigma c_1} \cap V \subset R^1_\sigma, \overrightarrow{b'_\sigma d_1} \cap V \subset R^2_\sigma$, and $\overrightarrow{a'_\sigma d_1} \cap V \subset R^3_\sigma$.*
3. *There are no stationary guard positions or guard segments different from r_i, r_j, r_l within the gadget from which any point of the corridor can be seen.*

Assuming that the above lemma holds, we will prove the following.

Lemma 22. *For any corridor to be attached at the right side of the polygon $\mathcal{P}$, the following properties are satisfied.*

1. *The intersection of any L-slab with any R-slab is contained in V.*

2. *For each* $\sigma \in \{0, \dots, 4n-1\}$*, it holds that* $\overrightarrow{a_\sigma c_0} \cap V \subset L^0_\sigma, \overrightarrow{b_\sigma c_0} \cap V \subset L^1_\sigma, \overrightarrow{a_\sigma d_0} \cap V \subset L^2_\sigma$*, and* $\overrightarrow{b_\sigma d_0} \cap V \subset L^3_\sigma$*.*

Proof. We will first show that the intersection of any L-slab with the line ℓ_c is contained in V. Consider Figure 19. Recall that o is the intersection point of the ray $\overrightarrow{a_0 c_0}$ with ℓ_c. The horizontal distance between ℓ_c and c_0 is 1. Let u be the intersection point of the lines ℓ_b and ℓ_r. From the polygon description in Section C.7 we know that $\|a_0 u\| = CN^2$, and $\|c_0 u\| \in [CN^2, CN^2 + 3k'] \subseteq [CN^2, CN^2 + 3N]$. The distance between the point $\frac{c_0+c_1}{2}$ and the point o is $\frac{\|c_0 u\|}{\|a_0 u\|} \in [1, 1 + \frac{3N}{CN^2}] = [1, 1 + 2N\rho]$. From the definition of the L-slabs, the intersection of the L-slabs with the vertical line ℓ_c is contained in the line segment $(o - (0, 2\varepsilon), o + (0, 4N\delta)) \subseteq (\frac{c_0+c_1}{2} + (0, 1 - 2\varepsilon), \frac{c_0+c_1}{2} + (0, 1 + 2N\rho + 4N\delta))$, which is contained in V as $\delta = 9\rho$.

By Property 1 of Lemma 21, any point in the intersection of an L-slab and an R-slab must have a y-coordinate within the range of V. As the angles of the slabs are exactly $\pi/4$ and $3\pi/4$, we get that also the x-coordinates of the intersection must be within the range of V, see also to the right of Figure 19. That gives us Property 1.

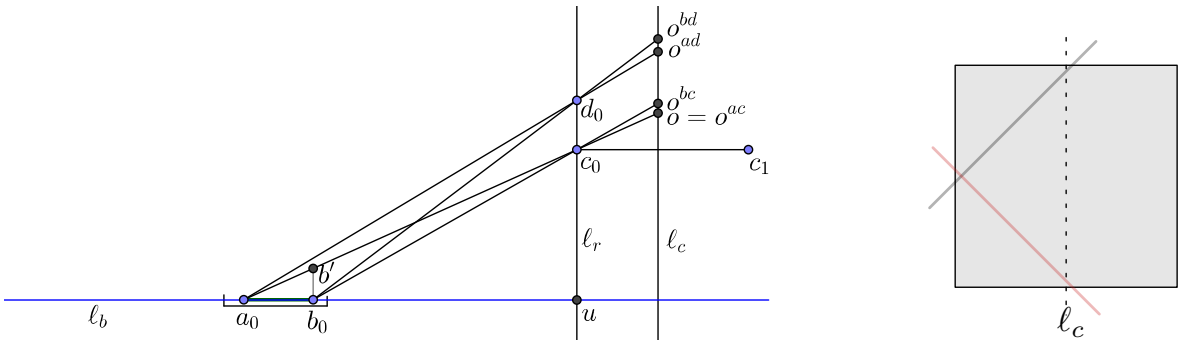

Figure 19: Left: Structure of rays with origins a_0 and b_0, containing points c_0 and d_0. Right: Even in the case that a left ray intersects ℓ_c at the very top of V and a right ray intersects at the very bottom of V, they still have to intersect within V.

For Property 2, let us first consider $\sigma = 0$. Let us define $o^{ac}, o^{ad}, o^{bc}, o^{bd}$ as the intersection points of the rays $\overrightarrow{a_0 c_0}, \overrightarrow{a_0 d_0}, \overrightarrow{b_0 c_0}, \overrightarrow{b_0 d_0}$ with the line ℓ_c (see Figure 19). We have $o = o^{ac}$. The points o^{ad}, o^{bc}, o^{bd} lie above o, and we will now estimate the distance between each of them and o. We do that as follows.

First, consider the distance $\|oo^{ad}\|$. From below, we have a trivial bound

$$\|oo^{ad}\| \geq \|c_0 d_0\| = \frac{3}{CN^2} = 2\rho.$$

From the similarity of triangles $a_0 c_0 d_0$ and $a_0 o o^{ad}$, and as the distance between the line ℓ_c and the line ℓ_r equals 1, we obtain the following upper bound for $\|oo^{ad}\|$

$$\|oo^{ad}\| = \|c_0 d_0\| \cdot \frac{\|a_0 u\| + 1}{\|a_0 u\|} = \frac{3}{CN^2} \cdot \frac{CN^2 + 1}{CN^2} = \frac{3}{CN^2} \cdot \left(1 + \frac{1}{CN^2}\right) \leq 2\rho + \frac{\varepsilon}{6}.$$

Let b' be a vertical projection of b_0 on the ray $\overrightarrow{a_0 c_0}$. From similarity of triangles $c_0 o o^{bc}$ and $c_0 b' b_0$, and triangles $a_0 b_0 b'$ and $a_0 u c_0$, we get the following equality

$$\|oo^{bc}\| = \|b_0 b'\| \cdot \frac{1}{\|b_0 u\|} = \|c_0 u\| \cdot \frac{\|a_0 b_0\|}{\|a_0 u\|} \cdot \frac{1}{\|b_0 u\|} = 3/2 \cdot \frac{\|c_0 u\|}{\|a_0 u\| \|b_0 u\|}.$$

We instantly get

$$\|oo^{bc}\| \geq 3/2 \cdot \frac{CN^2}{(CN^2)^2} = \rho.$$

For an upper bound, we compute

$$\frac{\|a_0u\|\|b_0u\|}{\|c_0u\|} \geq \frac{CN^2(CN^2-3/2)}{CN^2+3N} \geq \frac{(CN^2+3N)(CN^2-4N)}{CN^2+3N} = CN^2-4N \geq \frac{CN^2}{1+1/72},$$

where the last inequality follows since $C \geq 73 \cdot 4 = 292$. That gives us

$$\|oo^{bc}\| \leq 3/2 \cdot \frac{1+1/72}{CN^2} = \frac{3/2}{CN^2} + \frac{1/48}{CN^2} = \rho + \frac{\varepsilon}{6}.$$

In the same way as for $\|oo^{bc}\|$ we obtain the following bounds

$$\rho \leq \|o^{ad}o^{bd}\| \leq \rho + \varepsilon/6.$$

As $\|oo^{bd}\| = \|oo^{ad}\| + \|o^{ad}o^{bd}\|$, we instantly get

$$3\rho \leq \|oo^{bd}\| \leq 3\rho + \varepsilon/3.$$

Summarizing this part, we have the following bounds:

$$\rho \leq \|oo^{bc}\| \leq \rho + \varepsilon/3,\ 2\rho \leq \|oo^{ad}\| \leq 2\rho + \varepsilon/3,\ 3\rho \leq \|oo^{bd}\| \leq 3\rho + \varepsilon/3.$$

Therefore, the intersection points of the rays $\overrightarrow{a_0c_0}, \overrightarrow{b_0c_0}, \overrightarrow{a_0d_0}, \overrightarrow{b_0d_0}$ are contained in the required slabs, at a vertical distance of at most $\varepsilon/3$ from the centers of the slabs.

Now, consider any $\sigma \in \{0, \ldots, 4n-1\}$, $\tau \in \{a, b\}$ and $\eta \in \{c, d\}$. Let us denote by $\mathcal{R}$ all the rays with endpoint at one of the guard segments in the main area, going through c_0 and d_0 and all the rays from the endpoints of gadget guard segments through the points c_1, d_1.

Let $\widehat{o}$ be the intersection point of the ray $\overrightarrow{\tau_\sigma\eta_0}$ with the line ℓ_c (see Figure 20). We first

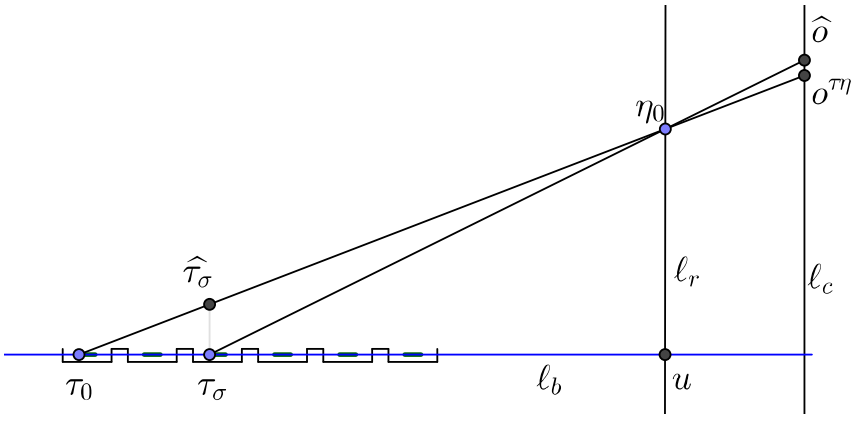

Figure 20: Bounding $\|o\widehat{o}\|$.

bound the distance $\|o^{\tau\eta}\widehat{o}\|$. Recall that u is the intersection point of ℓ_b and ℓ_r. Let $\widehat{\tau_\sigma}$ be the point on the ray $\overrightarrow{\tau_0\eta_0}$ vertically above τ_σ.

As the triangles $\eta_0 o^{\tau\eta}\widehat{o}$ and $\eta_0\widehat{\tau_\sigma}\tau_\sigma$ are similar, the triangles $\tau_0\tau_\sigma\widehat{\tau_\sigma}$ and $\tau_0 u\eta_0$ are similar, and the distance between the lines ℓ_c and ℓ_r is 1, we get the following equality

$$\|o^{\tau\eta}\widehat{o}\| = \|\tau_\sigma\widehat{\tau_\sigma}\| \cdot \frac{1}{\|\tau_\sigma u\|} = \|u\eta_0\| \cdot \frac{\|\tau_0\tau_\sigma\|}{\|\tau_0 u\|} \cdot \frac{1}{\|\tau_\sigma u\|} = 13.5\sigma \cdot \frac{\|u\eta_0\|}{\|\tau_0 u\|\|\tau_\sigma u\|}.$$

We first bound $\|o^{\tau\eta}\widehat{o}\|$ from above. As $C \geq 1297 \cdot 58 = 75226$, we get

$$\begin{aligned}\frac{\|\tau_0 u\|\|\tau_\sigma u\|}{\|u\eta_0\|} &\geq \frac{(CN^2 - 3/2)(CN^2 - 54N)}{CN^2 + 3N} \geq \frac{(CN^2 + 3N)(CN^2 - 58N)}{CN^2 + 3N} \\ &= CN^2 - 58N \geq \frac{CN^2}{1 + \frac{1}{1296N}},\end{aligned}$$

and, as $\sigma \leq 4N$, we can bound

$$\|o^{\tau\eta}\widehat{o}\| \leq 13.5\sigma \cdot \frac{1 + \frac{1}{1296N}}{CN^2} \leq \sigma \cdot \frac{13.5}{CN^2} + \frac{1/24}{CN^2} = \sigma\delta + \frac{\varepsilon}{3}.$$

To bound $\|o^{\tau\eta}\widehat{o}\|$ from below, we compute

$$\|o^{\tau\eta}\widehat{o}\| = 13.5\sigma \cdot \frac{\|u\eta_0\|}{\|\tau_0 u\|\|\tau_\sigma u\|} \geq 13.5\sigma \cdot \frac{CN^2}{(CN^2)^2} = \sigma\delta.$$

Therefore, as $\|o\widehat{o}\| = \|o^{\gamma\eta}\widehat{o}\| + \|oo^{\gamma\eta}\|$, the intersection points of the rays $\overrightarrow{a_\sigma c_0}, \overrightarrow{b_\sigma c_0}, \overrightarrow{a_\sigma d_0}, \overrightarrow{b_\sigma d_0}$ are contained in the required slabs, at a vertical distance of at most $2\varepsilon/3$ from the centers of the slabs.

We now need to verify that the rays stay in their respective slabs within the range of the x-coordinates of V. We therefore bound the slope of a ray $\overrightarrow{\tau_\sigma\eta_0}$ for any $\sigma \in \{0, \ldots, 4n-1\}$, $\tau \in \{a, b\}$ and $\eta \in \{c, d\}$. A bound from below is

$$\frac{\|u\eta_0\|}{\|\tau_\sigma u\|} \geq \frac{CN^2}{CN^2} \geq 1.$$

To bound the slope of $\overrightarrow{\tau_\sigma\eta_0}$ from above, we compute (as $C \geq 57 \cdot 1368 + 54 = 78030$)

$$\frac{\|u\eta_0\|}{\|\tau_\sigma u\|} \leq \frac{CN^2 + 3N}{CN^2 - 54N} \leq 1 + \frac{57N}{CN^2 - 54N} \leq 1 + \frac{1}{1368N}.$$

The center of each L-slab has the slope equal to 1. As the vertical distance between $\widehat{o}$ and the center of the slab is at most $2\varepsilon/3$, and $\frac{1}{1368N} \cdot 38N\rho = \frac{\rho}{36} = \frac{\varepsilon}{3}$, we get that $\overrightarrow{\tau_\sigma\eta_0} \cap V$ is contained in the corresponding slab. □

C.9.3 Constructing the corridor using slabs

We are now ready to describe the exact construction of the corridor. As mentioned before, the bottom wall is simply the line segment c_0c_1. We first describe the approximate shape of the upper wall, defined by a polygonal curve Λ connecting the points d_1 and d_0. Later we will present how to modify Λ into a final polygonal curve Λ', which is exactly the upper wall of the corridor.

Note that in the corridor construction here we assume that $i < j < l$. In particular, the L-slabs L_l^γ are above the L-slabs L_j^γ, which are above L_i^γ. For the R-slabs it is the opposite, i.e., the R-slabs R_l^γ are below the R-slabs R_j^γ, which are below R_i^γ.

Figure 21 shows the grid of slabs and a sketch of the curve Λ approximating the upper wall (excluding most of the leftmost and rightmost edge of Λ, with endpoints at d_0 and d_1,

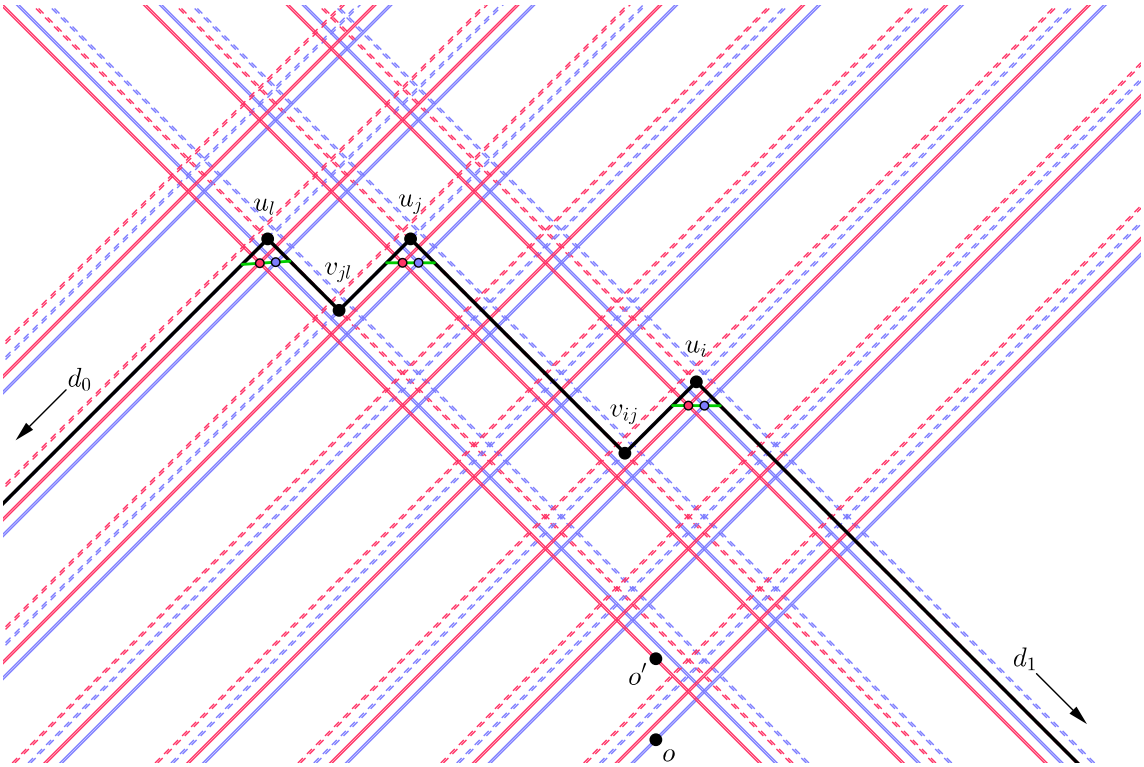

Figure 21: The grid of L- and R-slabs (bounded by blue and red lines) and an approximate shape Λ (black) of the upper wall of the corridor copying the guard segments s_i, s_j, s_l to the guard segments r_i, r_j, r_l in the gadget. Here, there are 8 guard segments $s_0, \ldots, s_7$ on the base line, and we have $i = 2$, $j = 5$, and $l = 6$. The blue lines bound the slabs corresponding to the rays originating at the left endpoints of the guard segments (i.e., slabs $L^0_\sigma, L^2_\sigma, R^1_\sigma, R^3_\sigma$), and the red lines bound the slabs corresponding to the rays originating at the right endpoints. The full lines bound the slabs corresponding to the rays passing through c_0 or c_1 (i.e., slabs $L^0_\sigma, L^1_\sigma, R^0_\sigma, R^1_\sigma$), and the dashed lines bound the slabs corresponding to the rays passing through d_0 or d_1. The blue points and the red points are the intersections of the rays $\overrightarrow{a_\sigma c_0} \cap \overrightarrow{a'_\sigma c_1}$ and $\overrightarrow{b_\sigma c_0} \cap \overrightarrow{b'_\sigma c_1}$, respectively, for $\sigma \in \{i, j, l\}$. The green segments each contains a critical segment (the part between the blue and the red point) of a copy-umbra for s_σ and r_σ with shadow corners c_0, c_1.

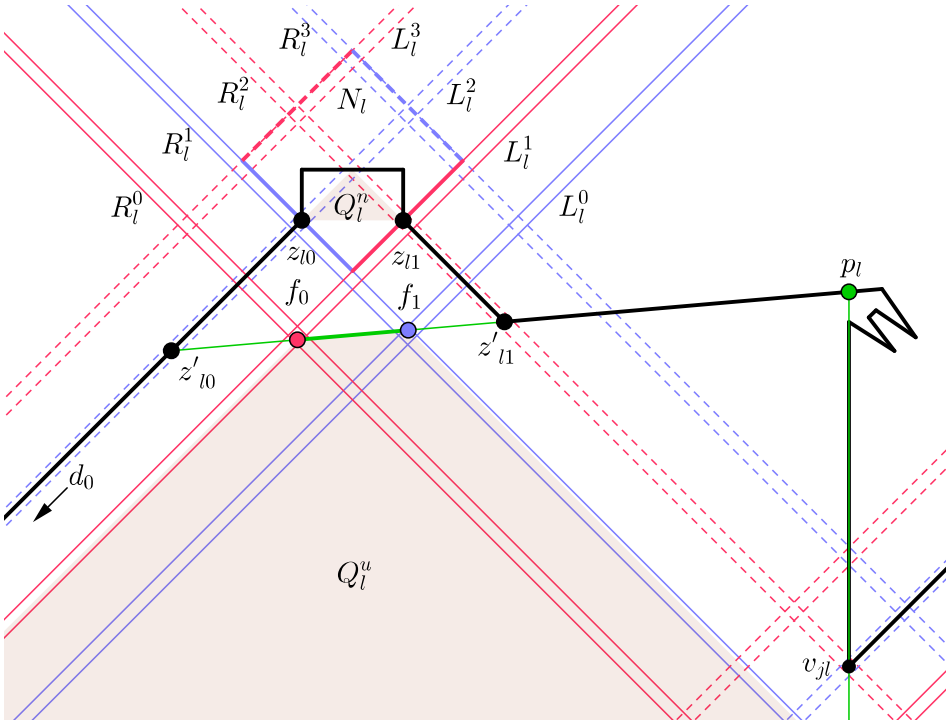

Figure 22: A closeup of the upper wall of the corridor from Figure 21 with features for copying the leftmost pair s_l, r_l of guard segments. The curve Λ' is drawn in black. The boundary of the square N_l is drawn with thick line segments. The points z_{l0}, z_{l1} are the shadow corners of the copy-nook Q_l^n (brown area) of s_l, r_l. The critical segment of Q_l^n is on the topmost black segment and it is so short that it appears as if it was just a point. The green line segment f_0f_1 is the critical segment of a copy-umbra Q_l^u of s_l, r_l with shadow corners c_0, c_1. The point p_l is a stationary guard position, from which a guard can see the area below the segment $z_{l0}z_{l1}$ containing f_0f_1. Furthermore, p_l sees the area to the left of the vertical ray emitting downwards from p_l.

respectively, since they are too long to be pictured together with the middle segments). For $\sigma \in \{i, j, l\}$, let u_σ be the intersection point of the rays $\overrightarrow{a_\sigma d_0}$ and $\overrightarrow{b'_\sigma d_1}$. Let v_{ij} be the intersection point of the rays $\overrightarrow{a_i d_0}$ and $\overrightarrow{b'_j d_1}$, and v_{jl} the intersection point of the rays $\overrightarrow{a_j d_0}$ and $\overrightarrow{b'_l d_1}$. The curve Λ is then a path defined by the points $d_1u_iv_{ij}u_jv_{jl}u_ld_0$. By Lemma 22, $\Lambda \cap V$ is contained in the union of the L-slabs and the R-slabs, as shown in Figure 21. Due to the relative position of the slabs $L_l^\gamma, L_j^\gamma, L_i^\gamma$ and $R_l^\gamma, R_j^\gamma, R_i^\gamma$ as discussed above, the curve Λ is x-monotone, and the point v_{ij} (resp. v_{jl}) has smaller y-coordinate than the neighbouring points u_i, u_j (resp. u_j, u_l), i.e., the curve Λ always has a zig-zag shape resembling the one from Figure 21.

We will now show how to modify Λ by adding to the curve some features. The first modification is in order to construct copy-nooks Q_i^n, Q_j^n, Q_l^n for each of the pairs (s_i, r_i), (s_j, r_j), and (s_l, r_l), respectively. Note that the area above c_0c_1 and below Λ already contains a copy-umbra Q_σ^u for each pair s_σ, r_σ for $\sigma \in \{i, j, l\}$ with shadow corners c_0 and c_1 (as Q_σ^u is contained in the triangular area bounded above c_0c_1 and below the rays $\overrightarrow{b_\sigma c_0}, \overrightarrow{a'_\sigma c_1}$, which, due to Lemmas 21 and 22, is below Λ). The second reason why we need to modify Λ is in order to create stationary guard positions p_i, p_j, p_l that see the areas of the copy-umbras, but nothing above their critical segments. In the following, we explain how to modify the fragment of Λ consisting of the leftmost two edges, i.e., the path $v_{jl}u_ld_0$. The construction is presented in Figure 22. We then perform similar modifications for the fragments of Λ consisting of the paths $v_{ij}u_jv_{jl}$ and $d_1u_iv_{ij}$.

First, we show how to construct a copy-nook Q_l^n of s_l and r_l with shadow corners at Λ. The curve Λ will then be modified so that Q_l^n is contained within the corridor. Let N_l be the square

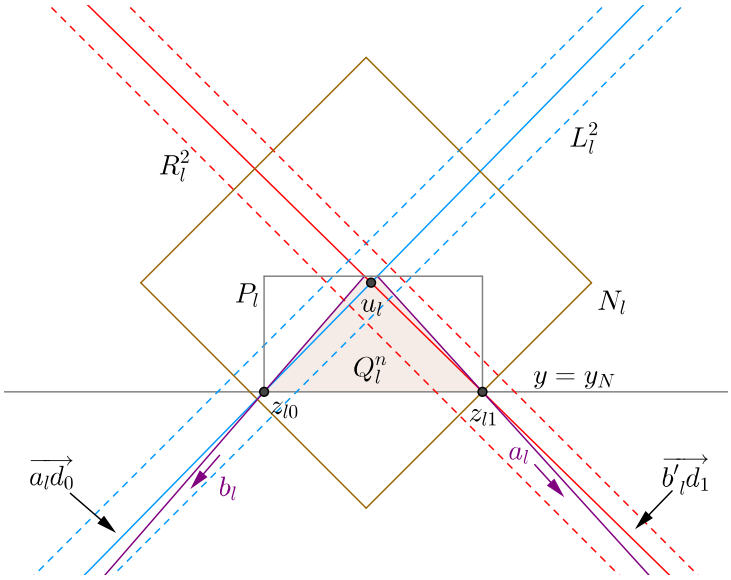

Figure 23: Construction of the nook Q_l^n.

consisting of points which are above the slabs L_l^1 and R_l^1, but not above L_l^3 or R_l^3. The square N_l is approximately the area which is seen both from the right endpoint b_l of s_l, and the left endpoint a'_l of r_l. Note that N_l contains the point u_l (as $u_l \in L_l^2 \cap R_l^2$, and thus u_l is above L_l^1 and R_l^1 and below L_l^3 and R_l^3). The copy-nook Q_l^n for the pair s_l, r_l will be created inside N_l (see Figure 23). Consider the two intersection points of the boundary of N_l with the line segments $v_{jl}u_l$ and $u_l d_0$. Let y_N be the larger of the y-coordinates of these two intersection points. The shadow corners of the nook Q_l^n are chosen as intersection points of the horizontal line $y = y_N$ with the line segments $v_{jl}u_l$ and $u_l d_0$, and they are denoted by z_{l1} and z_{l0}, respectively. In this way we ensure that both shadow corners are visible from any point within the segments s_l and r_l, and that they define a copy-nook Q_l^n for the pair of segments s_l, r_l. Note that the entire nook Q_l^n is contained in N_l, since by Lemmas 21 and 22 no point on s_l or r_l can see a point above the slabs L_l^3 or R_l^3. We now modify the curve Λ as follows. Let P_l be a quadrilateral with two vertices at z_{l0} and z_{l1}, and such that it contains vertical edges incident to z_{l0} and z_{l1}, and an edge containing the critical segment for Q_l^n. We modify Λ so that between the points z_{l0} and z_{l1}, it consists of the vertical edges and the topmost edge of P_l.

Now, consider the copy-umbra Q_l^u for the pair of segments s_l, r_l with shadow corners c_0 and c_1. Let $f_0 := \overrightarrow{b_l c_0} \cap \overrightarrow{b'_l c_1}$ and $f_1 := \overrightarrow{a_l c_0} \cap \overrightarrow{a'_l c_1}$. Note that the points f_0, f_1 correspond to the red and blue points in Figures 21 and 22. The segment $f_0 f_1$ is the critical segment for Q_l^u. By Lemmas 21 and 22, $f_0 \in L_l^1 \cap R_l^0$ and $f_1 \in L_l^0 \cap R_l^1$, and the squares $L_l^1 \cap R_l^0, L_l^0 \cap R_l^1$ have a sidelength of 2ε. Therefore the slope of the line $\overleftrightarrow{f_0 f_1}$ is in the interval $\left[-\frac{2\sqrt{2}\varepsilon}{\sqrt{2}\rho}, \frac{2\sqrt{2}\varepsilon}{\sqrt{2}\rho}\right] = [-1/6, 1/6]$, and this line intersects both line segments $v_{jl}u_l$ and $u_l d_0$. Let z'_{l0} and z'_{l1} be the intersection points of the line $\overleftrightarrow{f_0 f_1}$ with the line segments $u_l d_0$ and $u_l v_{jl}$, respectively. (We similarly define points z'_{i0}, z'_{i1} on $d_1 u_i v_{ij}$ as the intersection points with the line containing the critical segment of the umbra Q_i^u, and z'_{j0}, z'_{j1} on $v_{ij}u_j v_{jl}$ as the intersection points with the line containing the critical segment of the umbra Q_j^u.) We introduce a stationary guard position p_l by creating a pocket which will require modifying the curve Λ again. The pocket is extruding to the right from $v_{jl}u_l$, following the line $\overleftrightarrow{f_0 f_1}$, as pictured in Figure 22. Likewise, it is extruding vertically up from v_{jl}. The pocket contains a stationary guard position p_l on the line $\overleftrightarrow{f_0 f_1}$. Clearly, a guard placed at p_l sees nothing above the line segment $f_0 f_1$. Note that it sees the part of Q_l^u to the left of the

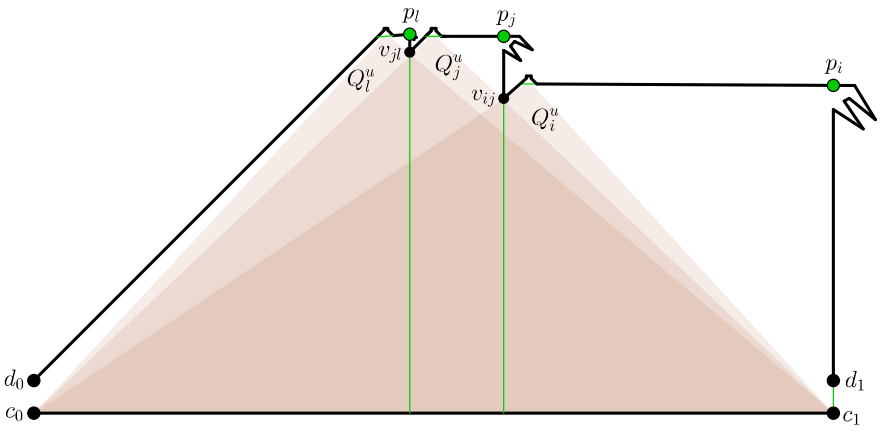

Figure 24: The complete construction of the corridor.

vertical line through v_{jl}.

For the middle two edges $v_{ij}u_jv_{jl}$ of Λ, we place the stationary guard position p_j vertically above v_{ij} so that it sees an area below the critical segment of the umbra Q_j^n and to the left of the vertical line through v_{ij}. For the rightmost edges $d_1u_iv_{jl}$, we place the stationary guard position p_i vertically above d_1 so that it sees an area below the critical segment of the umbra Q_i^u to the left of the vertical line through d_1. Let Λ' be the wall obtained by doing the modifications to Λ described here, and let C be the corridor, that is, the area bounded by the lower wall c_0c_1, the upper wall Λ' between d_0 and d_1, and by the vertical entrance segments c_0d_0 and c_1d_1. See Figure 24 for a picture of the complete corridor.

Lemma 23. *The stationary guard positions p_i, p_j, p_l have the following three properties.*

- *The three stationary guard positions p_i, p_j, p_l see together all of the corridor except the points above the segments $z'_{i0}z'_{i1}$, $z'_{j0}z'_{j1}$, $z'_{l0}z'_{l1}$.*
- *None of the guards can see anything to the right of the right entrance c_1d_1.*
- *None of the stationary guard positions p_i, p_j, p_l for the pairs $(s_i, r_i), (s_j, r_j), (s_l, r_l)$, respectively, can see any point on the critical segment of the nook or umbra of one of the other pairs.*

Proof. See Figure 24. For the first claim, note that the vertical lines through v_{jl} and v_{ij} divide the corridor into three parts. It is now clear that all points in the leftmost part below $z'_{l0}z'_{l1}$ are seen by p_l, all points in the middle part below $z'_{j0}z'_{j1}$ are seen by p_j, and all points in the rightmost part below $z'_{i0}z'_{i1}$ are seen by p_i.

For the second part, observe that the point p_i cannot see any point to the right of the vertical line through d_1, and the visibility of p_j and p_l is bounded by vertical lines more to the left.

For the last part, we note that the curve Λ' passes through points v_{jl} and v_{ij}, blocking visibility between stationary guard positions and critical segments corresponding to different pairs. □

Lemma 24. *Suppose that in each of the pairs (s_i, r_i), (s_j, r_j), (s_l, r_l) of guard segments corresponding to a corridor C, the two segments have the same orientation. Then C satisfies the following properties.*

1. *In any guard set G of $\mathcal{P}$ there are at least 3 guards placed within the corridor C, and if there are exactly 3 then they are placed at the stationary guard positions p_i, p_j, p_l. (The number is 2 instead of 3 if we construct C to copy only two segments.)*

2. *Let G be any set of points with exactly one guard on each guard segment and each stationary guard position, and with no guards outside of stationary guard positions and guard segments. If the whole area of C is seen by G, then for each of the pairs (s_i, r_i), (s_j, r_j), (s_l, r_l) the two guards on the segments specify the same value.*

3. *For any set of points G which satisfies the properties: (i) there is a guard at each point p_i, p_j, p_l and at each guard segment s_i, s_j, s_l and r_i, r_j, r_j, and (ii) the values specified by the pairs of segments s_i, s_j, s_l and r_i, r_j, r_j are consistent, G sees the whole area of C.*

4. *No guard at a stationary guard position or a guard segment outside the gadget can see any point inside the gadget below the line $\overleftrightarrow{d_0 c_1}$.*

5. *The vertical distance from c_0c_1 to the topmost point of the corridor is at most* 1.4.

Proof. For Property 1, note that the points defining the stationary guards within C can be seen only from within C. We can now use Lemma 14, setting A as the corridor area and choosing the points defining the stationary guards to construct the set M, to prove the desired property.

For Property 2, consider the set G as described, and let $\sigma \in \{i, j, l\}$. The stationary guard positions p_i, p_j, p_l cannot see any points above the line containing the critical segments of the umbra Q^u_σ. Lemma 23 and Property 3 of Lemma 21 give us that guards at the guard segments $s_0, s_1, \ldots, s_{4n-1}, r_i, r_j, r_l$ must see the critical lines of the nook and umbra Q^n_σ, Q^u_σ. We will now show that among all these segments, only guards placed at the segments s_σ, r_σ are able to do so. From Lemma 22 we get that for any $\sigma' > \sigma$, no guard on the guard segment $s_{\sigma'}$ can see a point in the square V below $L^0_{\sigma'}$. As $L^0_{\sigma'}$ is above L^3_σ, it is also above the critical segments of the nook Q^n_σ and the umbra Q^u_σ of the pair s_σ, r_σ. Likewise, for any $\sigma' < \sigma$, no guard on the guard segment $s_{\sigma'}$ can see a point in V above $L^3_{\sigma'}$. As $L^3_{\sigma'}$ is below L^0_σ, it is also below the critical segments of the pair s_σ, r_σ. A similar argument shows that among the guard segments r_i, r_j, r_l, only guards on r_σ can see any points on the critical segments for the pair s_σ, r_σ. Therefore the two guards on s_σ, r_σ must see together both critical segments of that pair, and by Lemma 17 the guards must specify the same value.

For Property 3, consider a set G satisfying (i) and (ii). By Lemma 23, the stationary guards can see the whole area of C except of the points which are above the line segments $z'_{i0}z'_{i1}, z'_{j0}z'_{j1}, z'_{l0}z'_{l1}$ containing the critical segments of the umbras. Consider any $\sigma \in \{i, j, l\}$, and the guard segments s_σ, r_σ. As these guards can see the complete critical segment for the umbra Q^u_σ, they can see the whole area contained above the critical segment and below Λ. As they can see the complete critical interval for the nook Q^n_σ, they can see the whole area of the polygon P_σ. Therefore, they can see the whole area above the critical interval and below Λ'.

Property 4 is a clear consequence of Lemma 23.

For Property 5, note that the top wall of the corridor can only extend beyond the square V due to a part of the wall Λ' creating a stationary guard position. Recall that the slope of the critical segments of the umbras Q^u_i, Q^u_j, Q^u_l is in the range $[-1/6, 1/6]$. Since $\|c_0c_1\| = 2$, it follows that any point on Λ' is at height at most $1 + 38N\rho + 2 \cdot 1/6 < 1.4$ above c_0c_1.

□

Lemma 25. *Assume that the endpoints of guard segments corresponding to a corridor C are at rational points, with the enumerators and the denominators upper-bounded by $(\zeta CN^2)^{O(1)}$. Then, we can construct the corridor C in such a way that each vertex of C has rational coordinates, with the enumerator and the denominator upper-bounded by $(\zeta CN^2)^{O(1)}$. The corridor construction can be done in polynomial time.*

Proof. Note that the entrances to the corridor are also at rational points, with enumerators and denominators upper-bounded by $(\zeta CN^2)^{O(1)}$. Therefore, each of the lines defining the polygonal curve Λ is defined by two rational points with this property. The same holds for the lines bounding the L-slabs and the R-slabs.

Consider the construction of a copy-nook Q^n_σ within the corridor. The vertices of Q^n_σ are then at points which are defined as intersection of two lines, where each line is defined by two rational points with enumerators and denominators upper-bounded by $(\zeta CN^2)^{O(1)}$. Therefore, the vertices of the nook, and therefore also the vertices of the quadrilateral P_σ are also of this form.

The stationary guard positions p_i, p_j, p_l are the intersection points of two lines, again each line defined by two points with the above property. Therefore, the enumerators and denominators of the stationary guard positions are also upper-bounded by $(\zeta CN^2)^{O(1)}$. The vertices of the pockets corresponding to p_i, p_j, p_l can be chosen with much freedom, and therefore they can also be at points satisfying the lemma statement. □

C.9.4 Corridor construction for gadgets at the left side of $\mathcal{P}$

For the gadgets attached at the left side of the polygon $\mathcal{P}$, the construction of the corridor is analogous. Now the points c_0, d_0 lie on the line ℓ_l instead of ℓ_r, and the points c_1, d_1 are to the left of c_0, d_0. As we want the points o (o') to correspond to the lowest intersection point of a ray from an endpoint of a guard segment in the base line (in the gadget, respectively) containing the point c_0 (c_1, respectively) with ℓ_c, we redefine these points in the following way. The point o is the intersection point of the ray $\overrightarrow{b_{4n-1}c_0}$ with ℓ_c, and o' is the intersection point of the ray $\overrightarrow{a'_l c_1}$ with ℓ_c. As we want the slabs L^γ_σ to contain fragments of rays from the endpoints of the segment s_σ, we redefine

$$L^\gamma_\sigma := S(o + (0, (4n - 1 - \sigma)\delta + \gamma\rho), \beta, \varepsilon).$$

Similarly, we redefine

$$R^\gamma_\sigma := S(o' + (0, \tau(\sigma)\delta + \gamma\rho), \alpha, \varepsilon).$$

As now the left endpoints of the gadget guard segments are further away from the line ℓ_c than the right endpoints, each gadget attached to the left side of $\mathcal{P}$ has to satisfy the following (instead of Lemma 21).

Lemma 26. *For any gadget to be attached to the left side of the polygon $\mathcal{P}$ and containing the guard segments $r_i = a'_i b'_i, r_j = a'_j b'_j, r_l = a'_l b'_l$ the following holds, where c_1 is the bottom-left endpoint of the corridor corresponding to the gadget.*

1. *The intersection of any R-slab with the line ℓ_c is contained in V.*
2. *For each $\sigma \in \{i, j, l\}$, it holds that $\overrightarrow{a'_\sigma c_1} \cap V \subset R^0_\sigma, \overrightarrow{b'_\sigma c_1} \cap V \subset R^1_\sigma, \overrightarrow{a'_\sigma d_1} \cap V \subset R^2_\sigma$, and $\overrightarrow{b'_\sigma d_1} \cap V \subset R^3_\sigma$.*
3. *There are no stationary guard positions or guard segments different from r_i, r_j, r_l within the gadget, from which any point of the corridor can be seen.*

For the same reason, instead of Lemma 22 we get the following lemma.

Lemma 27. *Within any corridor to be attached to the left side of the polygon $\mathcal{P}$ the following properties are satisfied.*

1. *The intersection of any L-slab with any R-slab is contained in V.*
2. *For each $\sigma \in \{0, \ldots, 4n-1\}$, it holds that $\overrightarrow{b_\sigma c_0} \cap V \subset L^0_\sigma, \overrightarrow{a_\sigma c_0} \cap V \subset L^1_\sigma, \overrightarrow{b_\sigma d_0} \cap V \subset L^2_\sigma$, and $\overrightarrow{a_\sigma d_0} \cap V \subset L^3_\sigma$.*

Due to symmetry of $\mathcal{P}_M$, the proof of Lemma 27 is the same as the proof of Lemma 22. The corridor construction and the proof of Lemma 24 is then the same as for the corridor attached at the right side of $\mathcal{P}$. The only difference is that in order to ensure that the L-slabs L_l^γ are above the L-slabs L_j^γ, which are above L_i^γ (which is required to get a meaningful zig-zag shape of the upper wall of the corridor) we now have to assume that $i > j > l$ (as the definition of the L-slabs has changed).

C.10 The $\geq$-addition gadget

In this section we present construction of the $\geq$-addition gadget which represents an inequality $x_i + x_j \geq x_l$, where $i, j, l \in \{0, \ldots, n-1\}$. In Section C.11 we will show how to modify the construction to obtain the $\leq$-addition gadget for the inequality $x_i + x_j \leq x_l$. For any equation of the form $x_i + x_j = x_l$ in Φ we will then add both gadgets into our polygon $\mathcal{P}$.

C.10.1 Idea behind the gadget construction

We first describe the general idea behind the construction of a gadget imposing an inequality $x'_i + x_j \geq x_l$ for three variables x'_i, x_j, x_l. See Figure 25 for a sketch of the construction. Let $w, v, h > 0$ be rational values such that $w > v + 3/2$. Let r'_i, r_j, r_l be right-oriented guard segments[‡] of length $3/2$ such that r_j has its left endpoint at the point $(-w, 0)$, r'_i has its right endpoint at $(w, 0)$, and r_l has its left endpoint at $(-2, -h)$. Let $g'_i := (w - 2 + x'_i, 0)$, $g_j := (-w - 1/2 + x_j, 0)$, and $g_l := (-5/2 + x_l, -h)$ be three guards on r'_i, r_j, r_l, respectively, representing the values $x'_i, x_j, x_l \in [1/2, 2]$.

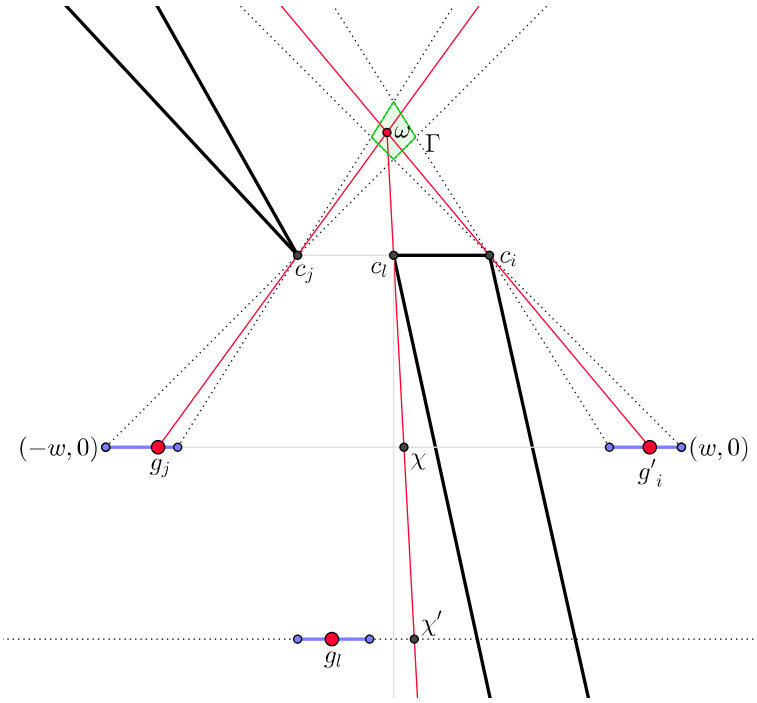

Figure 25: The thick black segments are edges of the polygon. The green quadrilateral is Γ. In order to see Γ together with g'_i and g_j, the guard g_l must be on or to the left of the point χ'.

Let $c_i := (v, h), c_j := (-v, h), c_l := (0, h)$. Let Γ be a collection of points ω such that the ray $\overrightarrow{\omega c_i}$ intersects r'_i, and the ray $\overrightarrow{\omega c_j}$ intersects r_j. Then Γ is a quadrilateral, bounded by the

[‡]We use r'_i instead of r_i here (and x'_i instead of x_i), as the guard segment r'_i specifying the value x'_i will only be a weak copy of the segment from the base line with a value x_i, i.e., it will hold that $x'_i \leq x_i$. More details will be provided later.

following rays: the rays with origin at the endpoints of r'_i and containing c_i, and the rays with origin at the endpoints of r_j and containing c_j. Suppose that

- for every point g'_i on r'_i and ω in Γ, the points ω and g'_i can see each other if and only if ω is on or to the right of the line $\overleftrightarrow{g'_i c_i}$,
- for every point g_j on r_j and ω in Γ, the points ω and g_j can see each other if and only if ω is on or to the right of the line $\overleftrightarrow{g_j c_j}$,
- for every point g_l on r_l and ω in Γ, the points ω and g_l can see each other if and only if ω is on or to the left of the line $\overleftrightarrow{g_l c_l}$.

Then, we can show the following result.

Lemma 28. *The guards g'_i, g_j, g_l can see together the whole quadrilateral Γ if and only if $x'_i + x_j \geq x_l$.*

Proof. Let $\omega \in \Gamma$ be the intersection point of the rays $\overrightarrow{g'_i c_i}$ and $\overrightarrow{g_j c_j}$.

Suppose that the guards g'_i, g_j, g_l together see the whole quadrilateral Γ. Since g'_i cannot see the area to the left of the line $\overleftrightarrow{\omega g'_i}$, and g_j cannot see the area to the left of the line $\overleftrightarrow{\omega g_j}$, there are points arbitrarily close to ω which are not seen by any of the guards g'_i, g_j. Therefore, g_l has to see ω.

Consider the rays $\overrightarrow{\omega c_i}$, $\overrightarrow{\omega c_j}$, and $\overrightarrow{\omega c_l}$. Let χ be the intersection point of the ray $\overrightarrow{\omega c_l}$ with a horizontal line $y = 0$, and χ' the intersection point of the ray $\overrightarrow{\omega c_l}$ with a horizontal line $y = -h$. Note that the guard g_l can see the point ω if and only if g_l is coincident with χ' or to the left of χ'.

From the similarity of triangles $g_j g'_i \omega$ and $c_j c_i \omega$ we get that $\frac{y(\omega)}{y(\omega)-h} = \frac{\|g'_i - g_j\|}{2v} = \frac{2w + x'_i - x_j - 3/2}{2v}$. From the similarity of triangles $g_j \chi \omega$ and $c_j c_\ell \omega$ we get that $\frac{y(\omega)}{y(\omega)-h} = \frac{\|\chi - g_j\|}{v}$, and therefore $\|\chi - g_j\| = w + x'_i/2 - x_j/2 - 3/4$, and $\chi = (x'_i/2 + x_j/2 - 5/4, 0)$. From the similarity of triangles $(0,0)\chi c_l$ and $(0,-h)\chi' c_l$ we get that $\chi' = (x'_i + x_j - 5/2, 0)$. The condition that the guard g_l is coincident with χ' or to the left of χ' is equivalent to $-5/2 + x_l \leq x'_i + x_j - 5/2$, i.e., $x'_i + x_j \geq x_l$.

On the other hand, if $x'_i + x_j \geq x_l$ then the guard g_l is coincident with χ' or to the left of χ', and therefore g_l can see ω. Then the guards g'_i, g_j, g_l can together see the whole Γ. □

C.10.2 Fragment of the gadget for testing the inequality

We now present construction of a polygon $\mathcal{P}_{\text{ineq}}$ containing three guard segments r'_i, r_j, r_l and vertices c_i, c_j, c_l with coordinates as described above, where we set $w := 26$, $v := 10$, and $h := 10.5$, enforcing an inequality on the values corresponding to the guard segments. The main part of the polygon $\mathcal{P}_{\text{ineq}}$ is pictured in Figure 26. The three blue line segments correspond to the guard segments, and the three green dots to stationary guard positions. The stationary guard positions g_t, g_m, g_b have been chosen as follows. The point g_t is at the ray with origin at the right endpoint of r_j and containing c_j, g_b is at the ray with origin at the right endpoint of r_j and containing c_l, and g_m is at the ray with origin at the right endpoint of r'_i and containing c_i. For each guard segment and each stationary guard position we introduce pockets of the polygon, such that the points defining each guard segment and stationary guard position cannot be seen from any other guard segment or stationary guard position. Two edges in the left of the figure are only shown partially. They end to the left at vertices $c_1 := (-CN^2, CN^2)$ and $d_1 := (-CN^2, CN^2 + 1.5)$, respectively. These vertices are connected by an edge $c_1 d_1$ which closes the polygon. We can show the following result.

Lemma 29. *Consider the polygon $\mathcal{P}_{ineq}$ from Figure 26. A set of points $G \subset \mathcal{P}_{ineq}$ of cardinality at most 6 guards $\mathcal{P}_{ineq}$ if and only if*

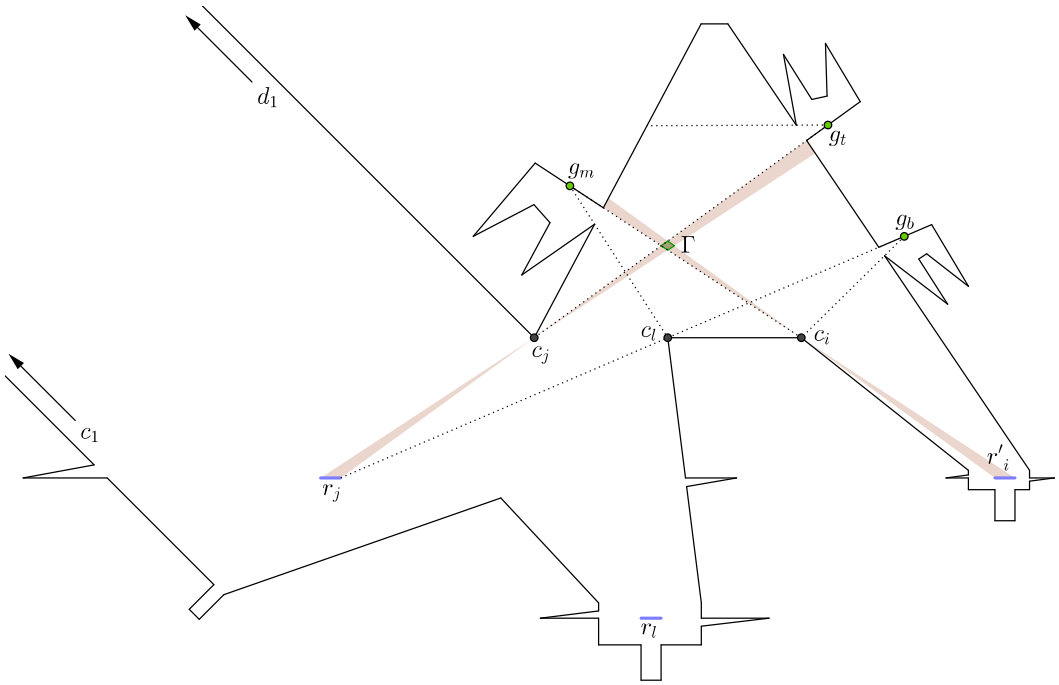

Figure 26: The main part of a polygon $\mathcal{P}_{\text{ineq}}$ for which all guard sets G of size 6 have the following form: (i) there is one guard at each stationary guard position (denoted by the green dots), (ii) there is one guard at each guard segment (blue line segments), and (iii) the values x'_i, x_j, x_l specified by the guards at the corresponding guard segments r'_i, r_j, r_l satisfy the inequality $x'_i + x_j \geq x_l$. Moreover, all sets G satisfying conditions (i)-(iii) are guard sets. The brown regions represent points of $\mathcal{P}_{\text{ineq}}$ included in some ray originating at r_j and containing c_j, or originating at r'_i and containing c_i. The quadrilateral Γ represents the intersection of these two regions.

- *there is exactly one guard placed at each guard segment r'_i, r_j, r_l and at each stationary guard position g_t, g_m, g_b, and*
- *the variables x'_i, x_j, x_l corresponding to the guard segments r'_i, r_j, r_l, respectively, satisfy the inequality $x'_i + x_j \geq x_l$.*

Proof. Assume first that the two conditions are satisfied. Observe that the stationary guard positions g_t, g_m, g_b have been chosen so that guards placed at them cannot see any point in the interior of Γ, but together with the guards g'_i, g_j, g_l placed at r'_i, r_j, r_l can see the whole area of $\mathcal{P}_{\text{ineq}} \setminus \Gamma$. (For this property to hold, the position of g'_i, g_j, g_l within the corresponding guard segments does not matter.) Since additionally the inequality $x'_i + x_j \geq x_l$ is satisfied, then Lemma 28 yields that the whole area of Γ is seen, and hence G guards $\mathcal{P}_{\text{ineq}}$.

Assume now that a set G of at most 6 guards sees all of $\mathcal{P}_{\text{ineq}}$. Using Lemmas 14 and 15 (where we set $A = \mathcal{P}_{\text{ineq}}$ and choose the points in M among the following: one point t_1 defining each stationary guard position, and one point q_2 defining each guard segment) we can show that the polygon requires at least 6 guards, and if there are 6 guards then there must be one guard at each stationary guard position, and at each guard segment. Then, as Γ is seen by the guards, by Lemma 28 we get that the inequality $x'_i + x_j \geq x_l$ holds. □

In the actual $\geq$-addition gadget we modify $\mathcal{P}_{\text{ineq}}$ in a way described later and scale it down by a factor of $\frac{1}{CN^2}$. Then we connect it to the main part of the polygon $\mathcal{P}$. The polygon $\mathcal{P}_{\text{ineq}}$ is attached at the right side of $\mathcal{P}$, and the connection between $\mathcal{P}_{\text{ineq}}$ and the main area of $\mathcal{P}$ is via a corridor. The point which corresponds to $(0, 0)$ in the polygon $\mathcal{P}_{\text{ineq}}$ is then at the position $m := c_1 + (1, -1)$ in $\mathcal{P}$, where c_1 is the bottom-right vertex of the corridor.

C.10.3 Copying guard segments to their final position

Let r'_i, r_j, r_l denote the guard segments of $\mathcal{P}_{\text{ineq}}$. We now show how to enforce dependency between appropriate guard segments from the base line of the polygon $\mathcal{P}$ and the guard segments r'_i, r_j, r_l.

We ensure that the guard segments r_j and r_l are copies of segments from the base line by connecting the (scaled) polygon $\mathcal{P}_{\text{ineq}}$ to the polygon $\mathcal{P}$ via a corridor. However, the segment r'_i cannot be copied in this way, as the edges of $\mathcal{P}_{\text{ineq}}$ are blocking visibility. Instead, we introduce an additional guard segment r_i within the gadget, and we copy appropriately chosen segments s_i, s_j, s_l from the base line into r_i, r_j, r_l. This part is explained in detail in Section C.10.5. Then, by introducing a copy-nook within the construction of $\mathcal{P}_{\text{ineq}}$, we ensure that r'_i is a weak copy of r_i. This is explained in detail in Section C.10.4. In Section C.10.6 we summarize the properties of the constructed gadget.

The gadget is scaled by a factor of $\frac{1}{CN^2}$ before it is attached to the corridor, and the points c_1, d_1 of the gadget are coincident with the points defining the right entrance of the corridor, which have the same names. After this operation, the middle point of the gadget (i.e., the point corresponding to $(0, 0)$ in the coordinate system of the gadget) satisfies the equality $m = c_1 + (1, -1)$, as stated in Section C.7.

For the picture of the complete gadget see Figure 27.

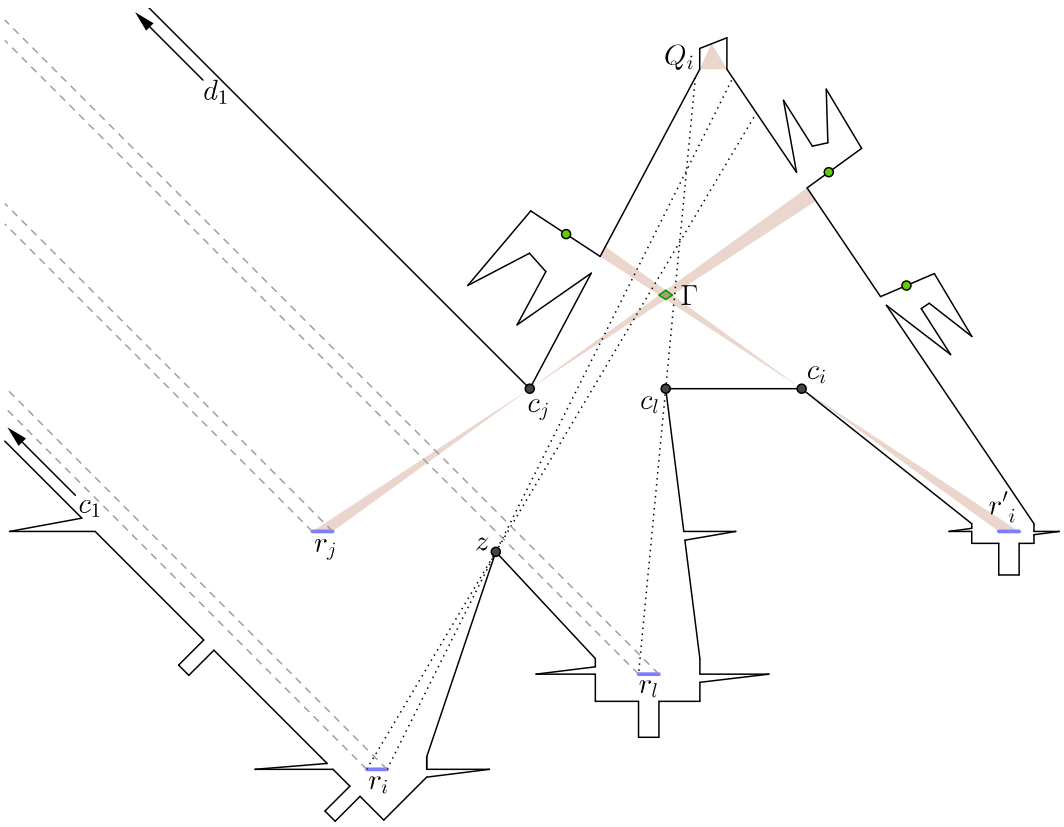

Figure 27: Detailed construction of the $\geq$-addition gadget. Note that the dotted lines show that no guard on r_i can see any point in Γ because of the corner z, a guard on r_i can always see both shadow corners of the copy-nook Q_i, and no point on r_l sees any point of Q_i because of the corner c_l. For each of the segments r_σ, $\sigma \in \{i, j, l\}$, the rays from points on r_σ through the corridor entrance $c_1 d_1$ are between the two grey dashed rays emitting from the endpoints of r_σ.

C.10.4 Introducing a new guard segment r_i

Consider the setting as described in Section C.10.2, and the polygon $\mathcal{P}_{\text{ineq}}$ from Figure 26. We explain how to modify $\mathcal{P}_{\text{ineq}}$ into a polygon $\mathcal{P}'_{\text{ineq}}$, which is a scaled version of our gadget. The main part of $\mathcal{P}'_{\text{ineq}}$ is shown in Figure 27, where again the edges to the left with endpoints at c_1 and d_1 are not fully shown.

The polygon $\mathcal{P}'_{\text{ineq}}$ is obtained from $\mathcal{P}_{\text{ineq}}$ in the following way. First, we add an additional guard segment r_i of length $3/2$ (i.e., the same as the length of other guard segments in the construction), with its left endpoint at the point $(-20.5, -17)$. This requires introducing a pocket corresponding to the added guard segment. We ensure that r'_i is a *weak copy* of r_i by creating a copy-nook Q_i the for pair of guard segments r_i, r'_i, which cannot be seen from any other guard segment or stationary guard position. The shadow corners of Q_i are $(2.5, 34)$ and $(4.5, 34)$. We have to ensure that after this modification, the gadget still enforces the desired inequality. In particular, we have to ensure that a guard placed at r_i cannot see any point in the interior of Γ. We can do that by introducing a new corner $z = (-12.5, -1.5)$ of the polygon that blocks r_i from seeing Γ. Note that z does not block r_i from seeing the segment c_1d_1 (as shown by the dashed grey lines in Figure 27).

Lemma 30. *A set of points $G \subset \mathcal{P}'_{ineq}$ of cardinality at most 7 guards $\mathcal{P}'_{ineq}$ if and only if*

- *there is exactly one guard placed at each guard segment r'_i, r_i, r_j, r_l and at each stationary guard position,*
- *the variables x_i, x'_i corresponding to the guard segments r_i, r'_i, respectively, satisfy the inequality $x_i \geq x'_i$, and*
- *the variables x'_i, x_j, x_l corresponding to the guard segments r'_i, r_j, r_l, respectively, satisfy the inequality $x'_i + x_j \geq x_l$.*

Proof. Assume that the polygon is guarded by a set G of at most 7 guards. Similarly as in Lemma 29 we can show that there must be exactly one guard at each guard segment and each stationary guard position. As the copy-nook Q_i can be seen only by guards placed at r_i and r'_i, it follows from Lemma 16 that the polygon is guarded by G only if the variables x_i, x'_i corresponding to r_i, r'_i, respectively, satisfy the inequality $x_i \geq x'_i$. As the quadrilateral Γ cannot be seen by a guard from r_i, we get from Lemma 29 that the variables x'_i, x_j, x_l corresponding to the guard segments r'_i, r_j, r_l, respectively, must satisfy the inequality $x'_i + x_j \geq x_l$.

Now assume that there is exactly one guard placed at each guard segment r'_i, r_i, r_j, r_l and at each stationary guard position, the variables x_i, x'_i satisfy the inequality $x_i \geq x'_i$, and the variables x'_i, x_j, x_l satisfy the inequality $x'_i + x_j \geq x_l$. Then the whole area of Γ is seen by the guards, as is the whole area of the nook Q_i. The remaining area is also seen by the guards, which we can show in the same way as in Lemma 29. □

C.10.5 Copying three guard segments via a corridor

We are given a formula from Φ of the form $x_i + x_j = x_l$, where $i, j, l \in \{0, \ldots, n-1\}$ and want to construct a gadget imposing an inequality $x_i + x_j \geq x_l$. We need to show that the values of the required variables can be copied into the guard segments r_i, r_j, r_l in the gadget described above. First, we will explain how to choose the segments from the base line to be copied into the gadget. Then, we will show that the $\geq$-addition gadget satisfies properties required by Lemma 21, which will ensure that the gadget can be connected to the main area by a corridor.

In order to apply a corridor construction as described in Section C.9 to copy three guard segments from the base line into the gadget, we require the segments in order from left to right on the base line to represent the variables x_i, x_j, x_l. Recall that there are n variables $x_0, \ldots, x_{n-1}$ in the formula Φ, but that for any $\sigma \in \{0, \ldots, n-1\}$, we use $x_{\sigma+n}, x_{\sigma+2n}$ and $x_{\sigma+3n}$ as synonyms

for x_σ. Therefore, the inequality $x_i + x_j \geq x_l$ is equivalent to $x_i + x_{n+j} \geq x_{2n+l}$. The guard segments on the base line are $s_0, \ldots, s_{4n-1}$, where each s_σ represents the variable x_σ. The segments $s_0, \ldots, s_{3n-1}$ are right-oriented whereas $s_{3n}, \ldots, s_{4n-1}$ are left-oriented. (In Section C.12 we explain how to obtain these dependencies between the guard segments.) Therefore, with slight abuse of notation, we redefine $j := j + n$ and $l := l + 2n$ so that $i < j < l < 3n$. Then, the guard segments s_i, s_j, s_l satisfy our requirements.

We now need to show that our gadget construction satisfies the conditions of Lemma 21. Recall that our gadget is the polygon $\mathcal{P}'_{\text{ineq}}$ scaled by a factor of $\frac{1}{CN^2}$. First, we will prove an auxiliary lemma.

Lemma 31. *Let G be any gadget to be attached at the right side of the polygon $\mathcal{P}$, and such that the guard segments r_i, r_j, r_l have a length of $\frac{3/2}{CN^2}$ and are contained in the box $m + [-\Delta, \Delta] \times [-\Delta, \Delta]$, where $\Delta := \frac{50}{CN^2}$. If for each $\sigma \in \{i, j, l\}$ the vertical distance from $\overrightarrow{b'_\sigma c_1} \cap \ell_c$ to the center of the slab R^0_σ is at most $\varepsilon/4$, then Property 2 of Lemma 21 holds for G.*

Proof. Fix any $\sigma \in \{i, j, l\}$. Recall that $m := c_1 + (1, -1)$, and that the vertical distance between the centers of two consecutive rays from $R^0_\sigma, R^1_\sigma, R^2_\sigma, R^3_\sigma$ is ρ. Therefore, it is enough if we show the following properties

1. the distance d^0_σ between the intersection points of $\overrightarrow{b'_\sigma c_1}$ and $\overrightarrow{b'_\sigma d_1}$ with ℓ_b is in the interval $[2\rho - \frac{\varepsilon}{4}, 2\rho + \frac{\varepsilon}{4}]$,
2. for any $\mu \in \{c, d\}$ the distance d^μ_σ between the intersection points of $\overrightarrow{a'_\sigma \mu_1}$ and $\overrightarrow{b'_\sigma \mu_1}$ with ℓ_b is in $\left[\rho - \frac{\varepsilon}{4}, \rho + \frac{\varepsilon}{4}\right]$, and
3. the absolute value of the slope of any ray with origin at an endpoint of r_σ and passing through a point in $c_1 d_1$ is in $\left[1 - \frac{1}{38N\rho} \cdot \frac{\varepsilon}{4}, 1 + \frac{1}{38N\rho} \cdot \frac{\varepsilon}{4}\right]$.

The first two properties yield that all rays $\overrightarrow{b'_\sigma c_1}, \overrightarrow{a'_\sigma c_1}, \overrightarrow{b'_\sigma d_1}$, and $\overrightarrow{a'_\sigma d_1}$ intersect the line ℓ_c within their corresponding slabs at the vertical distance of at most $\frac{3\varepsilon}{4}$ from the center of the slab. The last property yields that the rays are contained in the corresponding slabs throughout the square V. Assume in this proof without loss of generality that $c_1 = (0, 0)$.

We will prove Property 1 as follows. We have $|x(b'_\sigma) - x(m)| \leq \Delta \leq \frac{1}{1500}$ (as $C > 75000$), and

$$d^0_\sigma = \|c_1 d_1\| \cdot \frac{2 + x(b'_\sigma) - x(m)}{1 + x(b'_\sigma) - x(m)}$$
$$\in \left[\rho \cdot \frac{2 - \frac{1}{1500}}{1 + \frac{1}{1500}}, \rho \cdot \frac{2 + \frac{1}{1500}}{1 - \frac{1}{1500}}\right] = \left[\rho \cdot \left(2 - \frac{2}{1501}\right), \rho \cdot \left(2 + \frac{3}{1499}\right)\right] \subset \left[2\rho - \frac{\varepsilon}{4}, 2\rho + \frac{\varepsilon}{4}\right].$$

For Property 2, denote H as the distance between the point a'_σ and its vertical projection on the ray $\overrightarrow{b'_\sigma \mu_1}$. We have $\frac{d^\mu_\sigma}{H} = \frac{1}{1 + x(a'_\sigma) - x(m)} = \frac{1}{1 + x(b'_\sigma) - x(m) - \rho}$, $\frac{H}{\rho} = \frac{1 + y(m) - y(b'_\sigma) + \|\mu_1 c_1\|}{1 + x(b'_\sigma) - x(m)}$ and $\rho \leq \frac{1}{50000}$ (as $C \geq 75000$), and therefore

$$d^\mu_\sigma = \rho \cdot \frac{1 + y(m) - y(b'_\sigma) + \|\mu_1 c_1\|}{1 + x(b'_\sigma) - x(m)} \cdot \frac{1}{1 + x(b'_\sigma) - x(m) - \rho}$$
$$\in \left[\rho \cdot \frac{1 - \frac{1}{1500}}{1 + \frac{1}{1500}} \cdot \frac{1}{1 + \frac{1}{1500}}, \rho \cdot \frac{1 + \frac{1}{1500} + \frac{1}{50000}}{1 - \frac{1}{1500}} \cdot \frac{1}{1 - \frac{1}{1500} - \frac{1}{50000}}\right]$$
$$= \left[\rho \cdot \left(1 - \frac{4501}{2253001}\right), \rho \cdot \left(1 + \frac{458897}{224695603}\right)\right]$$
$$\subset \left[\delta - \frac{\varepsilon}{4}, \delta + \frac{\varepsilon}{4}\right].$$

For Property 3, we have to bound the slope of the rays. Note that $\frac{1}{38N\rho} \cdot \frac{\varepsilon}{4} = \frac{1}{1824N}$. Since $C \geq 50 \cdot 1824 = 91200$, we get that the slope is at least

$$\frac{y(c_1) - y(b'_\sigma)}{x(b'_\sigma) - x(c_1)} = \frac{1 + y(m) - y(b'_\sigma)}{1 + x(b'_\sigma) - x(m)} \geq \frac{1 - \Delta}{1 + \Delta} = 1 - \frac{50}{CN^2 + 50} \geq 1 - \frac{50}{CN} \geq 1 - \frac{1}{1824N}.$$

On the other hand, since $C \geq 101.5 \cdot 1824 + 50 = 185186$, we get that the slope is at most

$$\frac{y(d_1) - y(b'_\sigma)}{x(b'_\sigma) - x(d_1)} \leq \frac{1 + \Delta + \rho}{1 - \Delta} = 1 + \frac{101.5}{CN^2 - 50} \leq 1 + \frac{101.5}{(C - 50)N} \leq 1 + \frac{1}{1824N}.$$

□

Proof of Lemma 21 for the ≥-addition gadget. Note that in the ≥-addition gadget, the segments r_i, r_j, r_l have lengths of $\frac{3/2}{CN^2}$ and their right endpoints are placed at positions $m + (\frac{-20.5}{CN^2}, \frac{-17.5}{CN^2}), m + (\frac{-24.5}{CN^2}, 0), m + (\frac{-0.5}{CN^2}, \frac{-10.5}{CN^2})$, respectively, where $m := c_1 + (1, -1)$.

Recall that the point o' is the intersection point of the ray $\overrightarrow{b'_l c_1}$ with ℓ_c. The horizontal distance between ℓ_c and c_1 is 1. We have $b'_l := c_1 + (1, -1) + (\frac{-0.5}{CN^2}, \frac{-10.5}{CN^2})$. Then the vertical distance between c_0c_1 and the point o' is $\frac{CN^2+10.5}{CN^2-0.5} = 1 + \frac{11}{CN^2-0.5} \in [1, 1 + \frac{11.5}{CN^2}] \subset [1, 1 + 8\rho]$. The intersection of each of the R-slabs with the vertical line ℓ_c is therefore contained in an interval $[1 - 2\varepsilon, 1 + 8\rho + 3\delta] = [1 - 2\varepsilon, 1 + 35\rho]$ above c_0c_1, i.e., it is within the square V.

Let us denote the intersection points of the rays $\overrightarrow{b'_i c_1}$ and $\overrightarrow{b'_j c_1}$ by o'_i and o'_j, respectively. We have to bound the distances $\|o'o'_i\|$ and $\|o'o'_j\|$. We have

$$\|o'o'_j\| = y(o'_j) - y(o') = \frac{CN^2}{CN^2 - 24.5} - \frac{CN^2 + 10.5}{CN^2 - 0.5} = \frac{24.5}{CN^2 - 24.5} - \frac{11}{CN^2 - 0.5} \in \left[\delta - \frac{\varepsilon}{4}, \delta + \frac{\varepsilon}{4}\right],$$
$$\|o'o'_i\| = y(o'_i) - y(o') = \frac{CN^2 + 17.5}{CN^2 - 20.5} - \frac{CN^2 + 10.5}{CN^2 - 0.5} = \frac{38}{CN^2 - 20.5} - \frac{11}{CN^2 - 0.5} \in \left[2\delta - \frac{\varepsilon}{4}, 2\delta + \frac{\varepsilon}{4}\right].$$

We obtain Property 1. As the conditions of Lemma 31 are satisfied, Property 2 holds.

Property 3 also holds, as the edge c_jd_1 blocks all points at stationary guard positions and at the guard segment r'_i from seeing c_1d_1. □

C.10.6 Summary

Lemma 32. *Consider the addition gadget together with the corresponding corridor representing an inequality $x_i + x_j \geq x_l$, as described below. The following properties hold.*

- *The gadget and the corridor fit into a rectangular box of height* 3*.*
- *For any guard set of $\mathcal{P}$, at least* 10 *guards have to be placed in the corridor and the gadget.*
- *Assume that in the main area $\mathcal{P}_M$, there is exactly one guard at each guard segment, and there are no guards outside of the guard segments. Then* 10 *guards can be placed in the corridor and the gadget so that the whole corridor and gadget is seen if and only if the values x_i, x_j, x_l specified by the guards at the guard segments s_i, s_j, s_l satisfy the inequality $x_i + x_j \geq x_l$.*

Proof. Recall that by Lemma 24, the distance from c_0c_1 to the topmost point in the corridor is at most 1.4. The main part of the gadget is centered around the point $c_0 + (1, -1)$, and as it is of size $\Theta(\frac{1}{CN^2})$, the vertical space of at most 1.1 below the line segment c_0c_1 is enough to fit the gadget.

From Lemma 24, there are at least 3 guards placed within the corridor. From Lemma 30, there are at least 7 guards placed within the gadget. That gives us at least 10 guards needed.

Assume that there are exactly 10 guards within the corridor and gadget and that the corridor is completely seen. Then, from Lemma 24 and 30, there is exactly one guard at each guard segment and each stationary guard position. Then, by Lemma 24, the values x_i, x_j, x_l specified by s_i, s_j, s_l, and the values specified by r_i, r_j, r_l, are the same. By Lemma 30, the values x_i, x'_i corresponding to r'_i, r_i satisfy $x_i \geq x'_i$, and we also have $x'_i + x_j \geq x_l$. That enforces inequality $x_i + x_j \geq x_l$.

On the other hand, assume that $x_i + x_j \geq x_l$. We first place a guard at every of the 6 stationary guard positions in the corridor and gadget. By Lemmas 24 and 30, if we set guards at the 4 guard segments so that the values specified by guards at r_i, r_j, r_l are x_i, x_j, x_l, and the value x'_i specified by the guard at r'_i is the same as x_i, then the whole area of the gadget is guarded. □

C.11 The $\leq$-addition gadget

In this section we present construction of a gadget representing an inequality $x_i + x_j \leq x_l$, where $i, j, l \in \{0, \ldots, n-1\}$. The idea of the construction of this gadget is analogous as for the $\geq$-addition gadget presented in Section C.10, and the basic principle is presented in Section C.11.1. The principle underlying the $\geq$-addition gadget, as explained in Section C.10.1, required the polygon to have edges blocking the visibility between the segment r'_i (placed at the right side) and the segments r_j, r_l (placed at the left side). We managed to get around that by making r'_i a weak copy of a segment r_i, which in turn was a copy of a segment s_i on the base line. In contrast to this, the principle underlying the $\leq$-addition gadget presented here requires the polygon to have an edge separating r'_i placed at the left side from the segments r_j, r_l at the right side. If we were to build a gadget to be placed at the right side of $\mathcal{P}$, we would have to copy the variables corresponding to both of the segments r_j, r_l weakly, and the gadget would not enforce the desired inequality. To avoid this problem, we will place the gadget at the left side of $\mathcal{P}$. Then, we introduce an additional guard segment r_i, and we make r'_i a weak copy of r_i using a copy-nook Q_i. As r'_i is to the left of r_i, the copy-nook Q_i enforces the inequality $x'_i \geq x_i$, where r'_i, r_i represent x'_i, x_i, respectively. The result is that the gadget enforces the desired inequality. The relative placement of the segments r_i, r_j, r_l, r'_i (in particular, the value of w from Section C.10.2) has to be slightly different than in the construction of $\geq$-addition gadget, as it does not seem to be possible to make the gadgets completely symmetrical.

C.11.1 Idea behind the gadget construction

The idea behind a gadget imposing an inequality $x'_i + x_j \leq x_l$ is similar as for the $\geq$-addition gadget described above. As before, consider rational values $w, v, h > 0$, where $w > v + 3/2$, and let r'_i, r_j, r_l be right-oriented guard segments of length $3/2$ such that r'_i has its left endpoint at the point $(-w, 0)$, r_j has its right endpoint at $(w, 0)$, and r_l has its left endpoint at $(-2, -h)$. Let $g'_i := (-w - 1/2 + x_i, 0)$, $g_j := (w - 2 + x_j, 0)$, and $g_l := (-5/2 + x_l, -h)$ be three guards on r'_i, r_j, r_l, respectively, representing the values $x'_i, x_j, x_l \in [1/2, 2]$.

Suppose that there are vertices $c_i := (-v, h), c_j := (v, h), c_l := (0, h)$ of $\mathcal{P}$. As before, let Γ be a collection of points ω such that the ray $\overrightarrow{\omega c_i}$ intersects r'_i, and the ray $\overrightarrow{\omega c_j}$ intersects r_j. Then Γ is a quadrilateral, bounded by the following rays: the rays with origin at the endpoints of r'_i and containing c_i, and the rays with origin at the endpoints of r_j and containing c_j. Suppose that

- for every point g'_i on r'_i and ω in Γ, the points ω and g'_i can see each other if and only if ω is on or to the left of the line $\overleftrightarrow{g'_i c_i}$,
- for every point g_j on r_j and ω in Γ, the points ω and g_j can see each other if and only if ω is on or to the left of the line $\overleftrightarrow{g_j c_j}$,

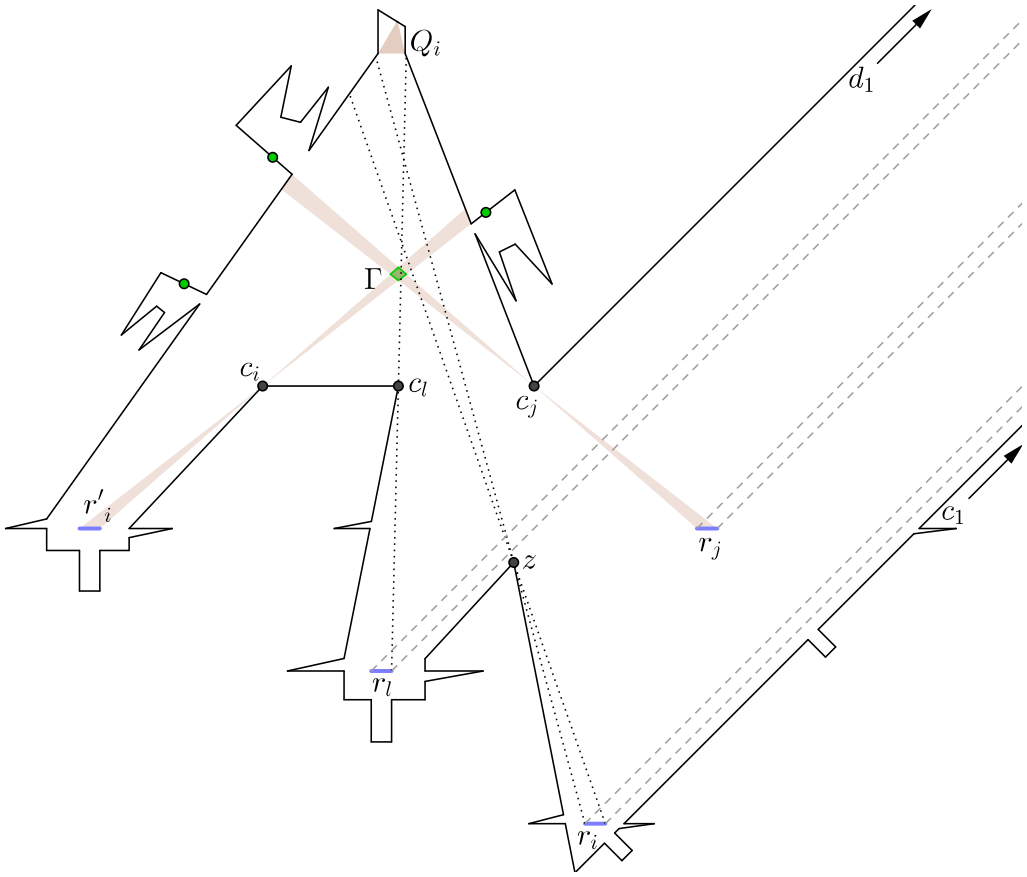

Figure 28: Detailed construction of the $\leq$-addition gadget. As previously, the dotted line shows that no guard on r_i can see any point in Γ because of the corner z, a guard on r_i can always see both shadow corners of the copy-nook Q_i, and no point on r_l sees any point of Q_i because of the corner c_l. For each of the segments r_σ, $\sigma \in \{i, j, l\}$, the rays from points on r_σ through the corridor entrance c_1d_1 are between the two grey dashed rays emitting from the endpoints of r_σ.

- for every point g_l on r_l and ω in Γ, the points ω and g_l can see each other if and only if ω is on or to the right of the line $\overleftrightarrow{g_l c_l}$.

Then, we can show the following result in an analogous way as we proved Lemma 28.

Lemma 33. *The guards g'_i, g_j, g_l can see together the whole quadrilateral Γ if and only if $x'_i + x_j \leq x_l$.*

C.11.2 A specification of the gadget

We will present the construction of a gadget with four guard segments r_i, r_j, r_l, r'_i, where the segments r'_i, r_j, r_l correspond to the segments in the idea described in Section C.11.1, where this time we set $w := 23.5$, $v := 10$, and $h := 10.5$. The gadget is shown in Figure 28 and should be attached to the left side of the main are $\mathcal{P}_M$ using a left corridor as described in Section C.9.4.

As for the case of the $\geq$-addition gadget, there are three stationary guards which do not see any point within Γ, but which enforce that the whole area except of Γ is seen whenever a guard is placed at each of the guard segments r_i, r_j, r_l, r'_i. The guard segment r_i has length $3/2$ (as do r'_i, r_j, r_l) and is placed with its left endpoint at the point $(13.75, -21.75)$. We will ensure that r'_i is a *weak copy* of r_i by creating a copy-nook Q_i the for pair of guard segments r_i, r'_i. The nook Q_i has shadow corners $(-1.5, 35)$ and $(0.5, 35)$. To ensure that a guard placed at r_i cannot see any point in the interior of Γ, we introduce a new corner $z := (8.5, -2.5)$ of the polygon that

blocks r_i from seeing Γ. Two edges to the right in the figure are not fully shown. They end at vertices $c_1 := (CN^2, CN^2)$ and $d_1 := (CN^2, CN^2 + 1.5)$.

In order to attach the gadget to the main are $\mathcal{P}_M$, we scale down the construction described here by the factor $\frac{1}{CN^2}$ and translate it so that the point which corresponds to $(0, 0)$ in the gadget will be placed at position $m := c_1 + (-1, -1)$. (Recall that the left entrance to the corridor, at which we attach the gadget, is the segment c_1d_1.)

The following lemma is proved in the same way as Lemma 30.

Lemma 34. *Let $\mathcal{P}'_{\text{rev-ineq}}$ be the polygon obtained from the gadget described above by closing it by adding the edge c_1d_1. A set of points $G \subset \mathcal{P}'_{\text{rev-ineq}}$ of cardinality at most 7 guards $\mathcal{P}'_{\text{rev-ineq}}$ if and only if*

- *there is exactly one guard placed at each guard segment r'_i, r_i, r_j, r_l and at each stationary guard position,*
- *the variables x_i, x'_i corresponding to the guard segments r_i, r'_i, respectively, satisfy the inequality $x_i \leq x'_i$, and*
- *the variables x'_i, x_j, x_l corresponding to the guard segments r'_i, r_j, r_l, respectively, satisfy the inequality $x'_i + x_j \leq x_l$.*

C.11.3 Copying three guard segments via a corridor

Here we proceed as in section C.10.5. We need to show is that the variables x_i, x_j, x_l can be copied into guard segments r_i, r_j, r_l from three guard segments on the base line. Due to the gadget construction, we now require the segment corresponding to the variable x_l to be the leftmost one, and all guard segments have to be right-oriented. With slight abuse of notation, we redefine $i := i + n$ and $j := j + 2n$. Then, the segments s_l, s_i, s_j satisfy our requirements.

As before, to prove that the corridor construction enforces the required dependency between the guards on the base line and guards within the gadget, i.e., for Lemma 24 to work, we need to show that our gadget construction satisfies the conditions of Lemma 26. This is done in much the same way as we did for the $\geq$-addition gadget. In particular, we use a symmetric version of Lemma 31 for gadgets attached to the left side of $\mathcal{P}$. The proofs are analogous, but with slightly different numbers.

C.11.4 Summary

In the same way as in Lemma 32, we get the following result.

Lemma 35. *Consider the $\leq$-addition gadget together with the corridor, corresponding to the inequality $x_i + x_j \leq x_l$. The following properties hold.*

- *The gadget and the corridor fit into a rectangular box of height 3.*
- *For any guard set of $\mathcal{P}$, at least 10 guards have to be placed in the corridor and the gadget.*
- *Assume that in the main area $\mathcal{P}_M$, there is exactly one guard at each guard segment, and there are no guards outside of the guard segments. Then 10 guards can be placed in the corridor and the gadget so that the whole corridor and gadget is seen if and only if the values x_i, x_j, x_l specified by the guards at the guard segments s_i, s_j, s_l satisfy the inequality $x_i + x_j \leq x_l$.*

C.12 The $\geq$- and $\leq$-orientation gadgets

In this section we explain how to enforce consistency between the guard segments on the base line which represent the same variable x_i, for $i \in \{0, \ldots, n-1\}$. Recall that there are four guard segments $s_i, s_{n+i}, s_{2n+i}, s_{3n+i}$ representing the variable x_i, and that the first three ones are right-oriented, and the last one is left-oriented.

We will present a gadget enforcing that two guard segments corresponding to the same variable x_i and oriented in different directions specify the variable consistently. We will then use this gadget for the following pairs of guard segments: (s_i, s_{3n+i}), (s_{n+i}, s_{3n+i}), and (s_{2n+i}, s_{3n+i}).

Consider two guard segments s_i, s_j on the base line, where s_i is right-oriented and s_j is left-oriented, and assume that there is one guard placed on each of these segments. Let x_i and x_j be the values represented by s_i and s_j, respectively. Let x_j^r be the value that would be specified by s_j if s_j was right-oriented instead of left-oriented. We have $x_j + x_j^r = 2.5$. Therefore s_i and s_j specify the same value if and only if $x_i + x_j^r = 2.5$.

Performing a simple modification of the $\geq$- and $\leq$-addition gadgets, we obtain the $\geq$- and $\leq$-orientation gadgets, which together enforce the equality $x_i + x_j^r = 2.5$. See Figure 29 for a detailed picture of the main part of the $\geq$-orientation gadget, which enforces that $x_i + x_j^r \geq 2.5$, or, equivalently, $x_i \geq x_j$. In the $\geq$-addition gadget, we copy three values from the base line into the gadget. Here, we copy only the value of the two segments s_i, s_j. Instead of the guard segment r_l inside the gadget, we create a stationary guard position p at the line containing r_l at distance $\frac{1}{CN^2}$ to the right of the right endpoint of r_l. Then p corresponds to the value of $5/2$ on r_l (ignoring that p lies outside r_l).

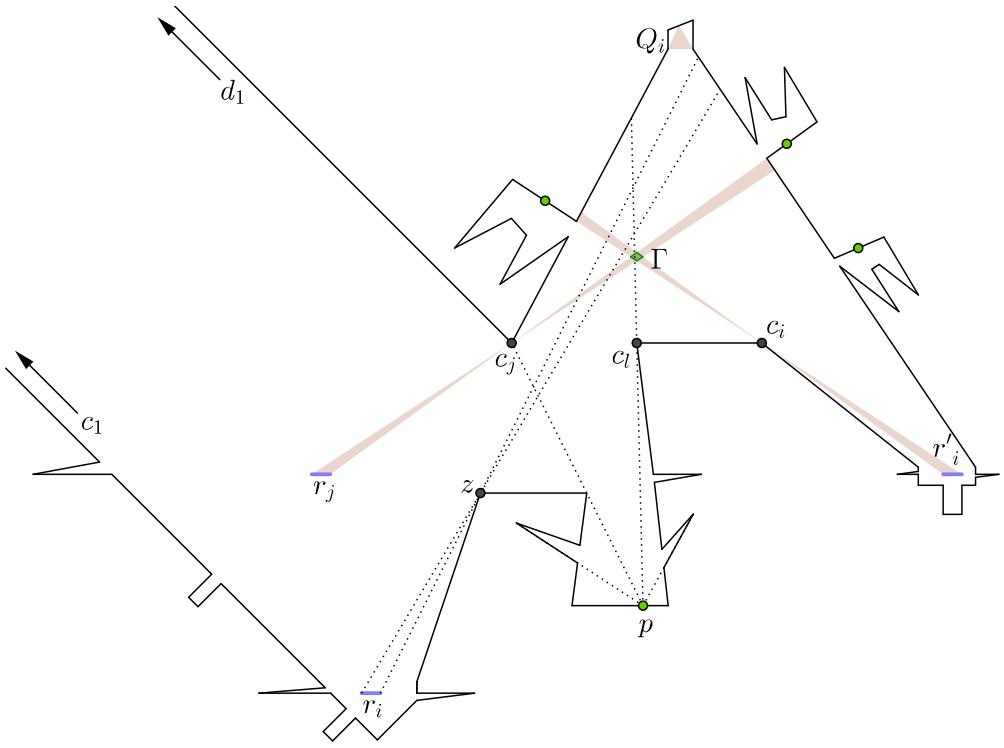

Figure 29: Detailed construction of the $\geq$-orientation gadget for $x_i^r + x_j^r \geq 2.5$, which is a modified version of the $\geq$-addition gadget.

The $\leq$-orientation gadget, which corresponds to the inequality $x_i + x_j^r \leq 5/2$, is obtained by an analogous modification of the $\leq$-addition gadget. Note that in both of these orientation gadgets, we create 4 stationary guard positions in the corridor instead of 6 for the addition gadgets, and the gadget itself contains 4 stationary guards and 3 guard segments, instead of 3 and 4 in the addition gadgets, respectively.

We summarize the properties of the orientation gadgets by the following Lemma, which can be proven in a way analogous to Lemma 30.

Lemma 36. *Consider the $\geq$-orientation gadget (resp. $\leq$-orientation gadget) together with the corresponding corridor for making r_i, r_j copies of guard segments s_i, s_j on the base line, where s_i is right-oriented and s_j is left-oriented. The following properties hold.*

- *The gadget and the corridor fit into a rectangular box of height 3.*
- *For any guard set of $\mathcal{P}$, at least 9 guards have to be placed in the corridor and the gadget.*
- *Assume that in the main area $\mathcal{P}_M$, there is exactly one guard at each guard segment, and there are no guards outside of the guard segments. Then 9 guards can be placed in the corridor and the gadget so that the whole corridor and gadget is seen if and only if the values x_i, x_j specified by the guards at the guard segments s_i, s_j satisfy the inequality $x_i \geq x_j$ (resp. $x_i \leq x_j$).*

C.13 The inversion gadget

In this section we present the construction of the inversion gadget which represent an inequality $x_i \cdot x_j = 1$, where $i, j \in \{0, \ldots, n-1\}$. We used Maple [25] for the construction and verification of this gadget.

C.13.1 Idea behind the gadget construction

We first describe the principle underlying the inversion gadget. See Figure 30. Let $a'_i := (1/2, 0)$, $b'_i := (2, 0)$, $a'_j := (13.9, 0.1)$, $b'_j := (15.4, 0.1)$, and suppose that $r_i := a'_i b'_i$ is a right-oriented guard segment representing the variable x_i and $r_j := a'_j b'_j$ is a left-oriented guard segment representing the variable x_j.

Let $\xi_0 := (7, \frac{541}{184}) \approx (7, 2.94)$ and $\xi_1 := (9, \frac{259139}{112792}) \approx (9, 2.30)$ and suppose that ξ_0, ξ_1 are shadow corners of an umbra Q_u with corners $\xi_0\xi_1 f_1 f_0$ of r_i, r_j. Then $f_0 := (\frac{499811}{70923}, \frac{38731813}{13049832}) \approx (7.05, 2.97)$ and $f_1 := (\frac{112379}{15432}, \frac{4355591}{1419744}) \approx (7.28, 3.07)$.

Lemma 37. *If guards p_i, p_j on r_i, r_j, respectively, see $f_0 f_1$ together, then $x_i x_j \leq 1$.*

Proof. Let π_i and π_j be the associated projections from r_i and r_j to $f_0 f_1$, respectively. Note that since p_i represents the variable x_i, we must have $p_i := (x_i, 0)$. Let

$$e := \pi_i(p_i) = \left(\frac{258288\, x_i - 16765}{36994\, x_i - 3065}, \frac{20013754\, x_i - 1295695}{6806896\, x_i - 563960} \right).$$

Now, $\pi_j^{-1}(e) = (15.9 - 1/x_i, 1/10)$, which represents the value $15.9 - (15.9 - 1/x_i) = 1/x_i$ on r_j. In order to see $f_0 f_1$ together with p_i, the guard p_j has to stand on $\pi_j^{-1}(e)$ or to the right. This corresponds to x_j being at most $1/x_i$. In other words, if a guard p_j on r_j sees $f_0 f_1$ together with p_i, then $x_i x_j \leq 1$. □

We now construct an umbra that impose the guards to satisfy the opposite inequality $x_i x_j \geq 1$: Let $\xi_2 := (7, \frac{8865}{752}) \approx (7, 11.79)$ and $\xi_3 := (9, \frac{4214815}{460976}) \approx (9, 9.14)$ and suppose that ξ_2, ξ_3 are shadow corners of a nook Q_n with corners $\xi_2\xi_3 f_3 f_2$ of r_i, r_j. Then $f_2 := (\frac{182083}{25835}, \frac{231222249}{19427920}) \approx (7.05, 11.90)$ and $f_3 := (\frac{205139}{28156}, \frac{130288905}{10586656}) \approx (7.29, 12.31)$.

Lemma 38. *If guards p_i, p_j on r_i, r_j, respectively, see $f_2 f_3$ together, then $x_i x_j \geq 1$.*

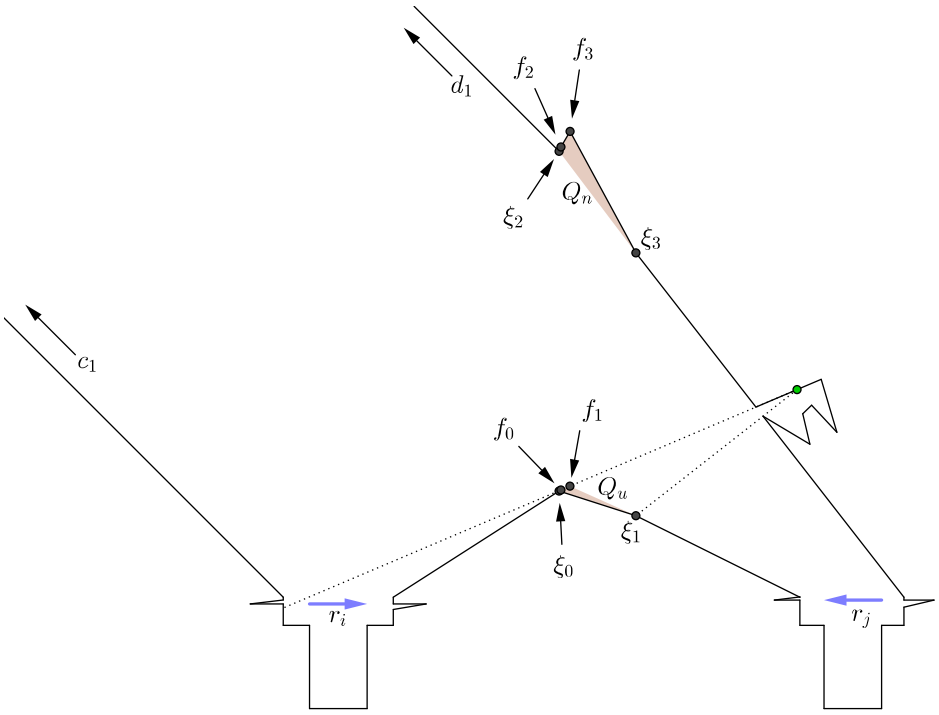

Figure 30: The inversion gadget. The nook and umbra (the brown areas) for the pair of guard segments r_i, r_j impose the inequality $x_i x_j = 1$ on the variables x_i and x_j represented by r_i and r_j. The green point is a stationary guard position which sees the umbra but nothing above the line $\overleftrightarrow{f_0 f_1}$.

Proof. Let $\hat{\pi}_0$ and $\hat{\pi}_1$ be the associated projections from r_i and r_j to $f_2 f_3$, respectively. Let

$$\hat{e} := \hat{\pi}_0(p_0) = \left(\frac{470184\, x_i - 29953}{67346\, x_i - 5517}, \frac{597022290\, x_i - 37933335}{50644192\, x_i - 4148784} \right).$$

Now, $\hat{\pi}_j^{-1}(\hat{e}) = (15.9 - 1/x_i, 1/10)$, which represents the value $15.9 - (15.9 - 1/x_i) = 1/x_i$ on r_j. In order to see $f_2 f_3$ together with p_i, the guard p_j has to stand on $\hat{\pi}_1^{-1}(\hat{e})$ or to the left. This corresponds to x_j being at least $1/x_i$ so that $x_i x_j \geq 1$. □

We thus have the following lemma:

Lemma 39. *If guards p_i and p_j placed at guard segments r_i and r_i, respectively, see both critical segments $f_0 f_1$ and $f_2 f_3$, then the corresponding values specified by p_i and p_j satisfy $x_i x_j = 1$.*

C.13.2 The construction of the gadget

We now explain how to make the complete gadget for the equation $x_i x_j = 1$, as shown in Figure 30. We make the wall in the gadget so that it creates the umbra Q_u and the nook Q_n as described before. We also create a stationary guard position at the green point in the figure which sees the umbra Q_u, but nothing above the line containing the critical segment of Q_u. Two edges at the left side of the gadget are not fully shown. They end at vertices $c_1 := (-CN^2, CN^2)$ and $d_1 := c_1 := (-CN^2, CN^2 + 1.5)$, respectively.

The gadget contains two guard segments r_i and r_j representing x_i and x_j, respectively, and it is required that r_i is right-oriented and r_j is left-oriented. Therefore, with slight abuse of notation, we redefine $j := j + 3n$, so that s_i, s_j are guard segments on the base line representing x_i, x_j, respectively, where s_i is right-oriented and s_j is left-oriented. We then use a corridor as described in Section C.9 to make r_i, r_j copies of s_i, s_j, respectively. Recall that the endpoints

of the right entrance of the corridor are denoted c_1, d_1. In order to attach the gadget to the corridor, we first scale it down by the factor $\frac{1}{CN^2}$ and then translate it so that the points c_1, d_1 of the gadget coincides with the endpoints of the right entrance of the corridor with the same names. We thus obtain that the point $(0,0)$ in the gadget becomes the point $m := c_1 + (1,-1)$ in $\mathcal{P}$.

Lemma 40. *Let $\mathcal{P}'_{inv}$ be the polygon obtained from the inversion gadget by closing it by adding the edge c_1d_1. A set of points $G \subset \mathcal{P}'_{inv}$ of cardinality at most 3 guards $\mathcal{P}'_{inv}$ if and only if*

- *there is exactly one guard placed at each guard segment r_i, r_j and at the stationary guard position,*
- *the variables x_i, x_j corresponding to the guard segments r_i, r_j, respectively, satisfy the equation $x_i \cdot x_j = 1$.*

Proof. Assume that the polygon is guarded by a set G of at most 3 guards. Similarly as in Lemma 29 we can show that there must be exactly one guard at each guard segment and at the stationary guard position. It then follows from Lemma 39 that $x_i \cdot x_j = 1$.

Now assume that there is exactly one guard placed at each guard segment r_i, r_j and at the stationary guard position, and that the variables x_i, x_j represented by the guards at r_i, r_j satisfy $x_i \cdot x_j = 1$. Then the whole area of Q_u and Q_n is seen by the guards. The remaining area is clearly also seen by the guards. □

C.13.3 Connecting the gadget with a corridor

We now need to show that our gadget construction satisfies the conditions of Lemma 21.

Proof of Lemma 21 for the inversion gadget. Note that in the inversion gadget, the segments r_i, r_j have lengths of $\frac{3/2}{CN^2}$ and their right endpoints are placed at positions $m + (\frac{2}{CN^2}, 0)$ and $m + (\frac{15.4}{CN^2}, \frac{0.1}{CN^2})$, respectively, where $m := c_1 + (1,-1)$.

Recall that the point o' is the intersection point of the ray $\overrightarrow{b'_j c_1}$ with ℓ_c. The horizontal distance between ℓ_c and c_1 is 1. We have $b'_j := c_1 + (1,-1) + (\frac{15.4}{CN^2}, \frac{0.1}{CN^2})$. Then the vertical distance between c_0c_1 and the point o' is $\frac{CN^2-0.1}{CN^2+15.4} = 1 - \frac{15.5}{CN^2+15.4} \in [1 - \frac{15.5}{CN^2+15.4}, 1] \subset [1-11\rho, 1]$. The intersection of each of the R-slabs with the vertical line ℓ_c is therefore contained in an interval $[1 - 11\rho - 2\varepsilon, 1 - 11\rho + 3\delta] = [1 - 11\rho - 2\varepsilon, 1 + 25\rho]$ above c_0c_1, i.e., it is within the square V. We obtain Property 1.

Let us denote the intersection point of the ray $\overrightarrow{b'_i c_1}$ and ℓ_c by o'_i. We have to bound the distance $\|o'o'_i\|$. We have

$$\|o'o'_i\| = y(o'_i) - y(o') = \frac{CN^2}{CN^2+2} - \frac{CN^2-0.1}{CN^2+15.4} = \frac{-2}{CN^2+2} - \frac{-15.5}{CN^2+15.4} \in \left[\delta - \frac{\varepsilon}{4}, \delta + \frac{\varepsilon}{4}\right].$$

As the conditions of Lemma 31 are satisfied, Property 2 holds.

Property 3 also holds, as the stationary guard position in the gadget cannot see the edge c_1d_1. □

C.13.4 Summary

Lemma 41. *Consider the inversion gadget together with the corresponding corridor representing an inequality $x_i \cdot x_j = 1$, as described below. The following properties hold.*

- *The gadget and the corridor fit into a rectangular box of height 3.*
- *For any guard set of $\mathcal{P}$, at least 5 guards have to be placed in the corridor and the gadget in total.*

- *Assume that in the main area $\mathcal{P}_M$, there is exactly one guard at each guard segment, and there are no guards outside of the guard segments. Then the whole area of the corridor and the gadget can be guarded by 5 guards (together with the guards on the base line) if and only if the values x_i, x_j specified by the guards at the guard segments s_i, s_j satisfy the inequality $x_i \cdot x_j = 1$.*

Proof. The proof for the first property is similar to that for the addition gadget in Lemma 32.

From Lemma 24, there are at least 2 guards placed within the corridor. Furthermore, there must be at least 3 guards placed within the gadget by Lemma 40. That gives us at least 5 guards needed.

Assume that there are exactly 5 guards within the corridor and gadget and that the corridor is completely seen. Then, from Lemma 24 and 40, there is exactly one guard at each guard segment and each stationary guard position. Then, by Lemma 24, the values x_i, x_j specified by s_i, s_j, and the values specified by r_i, r_j, are the same. Lemma 24 gives that the guards at r_i, r_j must see the critical segments of both Q_n and Q_u. Then, by Lemma 40, the values x_i, x_j thus satisfy $x_i \cdot x_j = 1$.

On the other hand, assume that $x_i \cdot x_j = 1$. We first place a guard at every of the 3 stationary guard positions in the corridor and gadget. By Lemmas 24 and 40, if we set guards at the 2 guard segments so that the values specified by guards at r_i, r_j are x_i, x_j, then the whole area of the gadget is guarded. □

C.14 Putting it all together

Let Φ be an ETR-INV formula with n variables, k_1 equations of the form $x_i + x_j = x_l$, and k_2 equations of the form $x_i \cdot x_j = 1$. We have already explained how to construct the polygon $\mathcal{P}(\Phi)$, but we shall here give a brief summary of the process. We start by constructing the main area with $4n$ guard segments. We modify the pockets corresponding to the variables x_i for which Φ contains equation $x_i = 1$, as described in Section C.8. To enforce dependency between the base line guard segments corresponding to the same variable, we construct $3n$ $\geq$-orientation gadgets (attached at the right side of the polygon) and $3n$ $\leq$-orientation gadgets (attached at the left side), as described in Section C.12. For each equality of the form $x_i + x_j = x_l$ in Φ, we construct a corresponding $\geq$-addition gadget (attached at the right side), and a $\leq$-addition gadget (attached at the left side), as described in Sections C.10 and C.11, respectively. For each equality of the form $x_i \cdot x_j = 1$ in Φ we construct a corresponding inversion gadget (attached at the right side), as described in Section C.13. The total number of gadgets at each side of $\mathcal{P}$ is therefore at most $3n + k_1 + k_2 = k'$, as stated in Section C.7.

Without loss of generality, we assume that the y-coordinate of the base line of $\mathcal{P}(\Phi)$ is 0. We set $g(\Phi) := 58n + 20k_1 + 5k_2$. Then $(\mathcal{P}(\Phi), g(\Phi))$ is an instance of the art gallery problem, and the following theorem holds.

Theorem 42. *The polygon $\mathcal{P}(\Phi)$ has vertices at rational coordinates, which can be computed in polynomial time. Moreover, there exist constants $d_0, \ldots, d_{n-1} \in \mathbb{Q}$ such that for any $x := (x_0, \ldots, x_{n-1}) \in \mathbb{R}^n$, x is a solution to Φ if and only if there exists a guard set G of cardinality $g(\Phi)$ containing guards at all the positions $(x_0 + d_0, 0) \ldots, (x_{n-1} + d_{n-1}, 0)$.*

Proof. Consider a guard set G of the polygon $\mathcal{P} := \mathcal{P}(\Phi)$. By Lemma 20, G has at least $4n$ guards placed in $\mathcal{P}_M$, and if the number of guards within $\mathcal{P}_M$ equals $4n$, then there must be exactly one guard at each guard segment. Lemma 36 implies that within each of the $6n$ orientation gadgets of $\mathcal{P}$ together with the corresponding corridors, there are at least 9 guards, giving at least $54n$ guards in total. Similarly, from Lemmas 32 and 35 we obtain that there must be at least 10 guards placed within each $\geq$-addition gadget and each $\leq$-addition gadget plus the corresponding corridors, giving at least $20k_1$ guards. By Lemma 41, there are at least 5

guards within each inversion gadget and the corresponding corridor, giving at least $5k_2$ guards in total. Therefore, G has at least $58n + 20k_1 + 5k_2$ guards, which is equal to $g(\Phi)$.

If a guard set of size $g(\Phi)$ exists, then there are exactly $4n$ guards in $\mathcal{P}_M$, 9 guards within each orientation gadget, 10 guards within each $\geq$-addition gadget and each $\leq$-addition gadget, and 5 guards within each inversion gadget. The same lemmas give us then that there is exactly one guard at each guard segment and each stationary guard position, and no guards away from the guard segments or the stationary guard positions. Also, the variables $x_0, \ldots, x_{n-1}$ specified by the guard segments $s_0, \ldots, s_{n-1}$ are a solution to Φ.

On the other hand, if there exists a solution to Φ, then we get a guard set of size $g(\Phi)$ by placing the guards accordingly. It is thus clear that the solutions to Φ correspond to the optimal guard sets of $\mathcal{P}$, as stated in the theorem.

Due to Lemmas 19 and 25, we get that the vertices of $\mathcal{P}_M$ and the corridor vertices are all rational, with the enumerators and denominators polynomially bounded. Next, consider all the vertices of the gadgets. Each gadget is first described as a polygon with coordinates where enumerators and denominators are both of size $\Theta(1)$ (as this construction is fixed and it does not depend on the formula Φ; and we can easily choose the vertices so that they are all at rational coordinates), and this polygon is subsequently scaled down by a factor of $\frac{1}{CN^2}$ and attached at a corridor entrance, which also has polynomially bounded enumerators and denominators. Thus, the coordinates of the vertices in the gadgets have polynomially bounded complexity and can be computed in polynomial time. □

We can now prove the main theorem of the paper.

Theorem 1. *The art gallery problem is $\exists\mathbb{R}$-complete, even the restricted variant where we are given a polygon with vertices at integer coordinates.*

Proof. By Theorem 4, the art gallery problem is in the complexity class $\exists\mathbb{R}$. From Theorem 6 we know that the problem ETR-INV is $\exists\mathbb{R}$-complete. We presented a polynomial time construction of an instance $(\mathcal{P}(\Phi), g(\Phi))$ of the art gallery problem from an instance Φ of ETR. Theorem 42 gives that it is $\exists\mathbb{R}$-hard to solve the art gallery problem when the coordinates of the polygon vertices are given by rational numbers. Note that the number of vertices of $\mathcal{P}$ is proportional to the input length $|\Phi|$. By Theorem 42, there is a polynomial $|\Phi|^m$ which is a bound on every denominator of a coordinate of a vertex in $\mathcal{P}$. The product Π of denominators of all coordinates of vertices of $\mathcal{P}$ thus has size at most $|\Phi|^{m\,O(|\Phi|)}$. It follows that we can express Π by $O(m|\Phi|\log|\Phi|)$ bits. By multiplying every coordinate of $\mathcal{P}$ by Π, we get a polygon $\mathcal{P}'$ with integer coordinates and the theorem follows. □

Theorem 2, restated below, likewise easily follows from Lemma 13 in Appendix B and Theorem 42.

Theorem 2. *Let Φ be an ETR formula with k variables. Then there is an instance $(\mathcal{P}, g)$ of the art gallery problem, and constants $c_1, d_1, \ldots, c_k, d_k \in \mathbb{Q}$, such that*

- *if Φ has a solution, then $\mathcal{P}$ has a guard set of size g, and*
- *for any guard set G of $\mathcal{P}$ of size g, there exists $(x_1, \ldots, x_k) \in S_\Phi$ such that G contains guards at positions $(c_1x_1 + d_1, 0), \ldots, (c_kx_k + d_k, 0)$.*

We can now prove Corollary 3, restated below.

Corollary 3. *Given any real algebraic number α, there exists a polygon $\mathcal{P}$ with vertices at rational coordinates such that in any optimal guard set of $\mathcal{P}$ there is a guard with an x-coordinate equal to α.*

Proof. Let $P(x)$ be a polynomial of degree more than 0 in one variable x such that the equation $P(x) = 0$ has α as a solution. The equation might have other solutions as well, but we can choose integers p_1, p_2, q_1, q_2 such that α is the only solution in the interval $[p_1/q_1, p_2/q_2]$. Then the formula $P(x) = 0 \wedge p_1 \leq q_1 x \wedge q_2 x \leq p_2$ is an instance of the problem ETR with a unique solution $x = \alpha$. Now, by Theorem 2, there exists a polygon $\mathcal{P}$ and rational constants c, d such that in any optimal guard set of $\mathcal{P}$, one guard has coordinates $(c\alpha + d, 0)$. By subtracting d from the x-coordinate of all vertices of $\mathcal{P}$ and then dividing all coordinates by c, we get a polygon $\mathcal{P}'$ such that any optimal guard set of $\mathcal{P}'$ has a guard at the point $(\alpha, 0)$. □

References

[1] Zachary Abel, Erik D. Demaine, Martin L. Demaine, Sarah Eisenstat, Jayson Lynch, and Tao B. Schardl. Who needs crossings? Hardness of plane graph rigidity. In *32nd International Symposium on Computational Geometry (SoCG 2016)*, pages 3:1–3:15, 2016.

[2] Mikkel Abrahamsen, Anna Adamaszek, and Tillmann Miltzow. Irrational guards are sometimes needed. In *33rd International Symposium on Computational Geometry (SoCG 2017)*, 2017. To appear.

[3] Alok Aggarwal. *The art gallery theorem: its variations, applications and algorithmic aspects*. PhD thesis, 1984.

[4] Saugata Basu, Richard Pollack, and Marie-Françoise Roy. On the combinatorial and algebraic complexity of quantifier elimination. *Journal of the ACM*, 43(6):1002–1045, 1996.

[5] Patrice Belleville. Computing two-covers of simple polygons. Master's thesis, McGill University, 1991.

[6] Daniel Bienstock. Some provably hard crossing number problems. *Discrete & Computational Geometry*, 6(3):443–459, 1991.

[7] Édouard Bonnet and Tillmann Miltzow. Parameterized hardness of art gallery problems. In *24th Annual European Symposium on Algorithms (ESA)*, pages 19:1–19:17, 2016.

[8] Dorit Borrmann, Pedro J. de Rezende, Cid C. de Souza, Sándor P. Fekete, Stephan Friedrichs, Alexander Kröller, Andreas Nüchter, Christiane Schmidt, and Davi C. Tozoni. Point guards and point clouds: Solving general art gallery problems. In *Symposuim on Computational Geometry 2013 (SoCG 2013)*, pages 347–348, 2013.

[9] Björn Brodén, Mikael Hammar, and Bengt J. Nilsson. Guarding lines and 2-link polygons is APX-hard. In *Proceedings of the 13th Canadian Conference on Computational Geometry (CCCG 2001)*, pages 45–48, 2001.

[10] John Canny. Some algebraic and geometric computations in PSPACE. In *Proceedings of the twentieth annual ACM symposium on Theory of computing (STOC 1988)*, pages 460–467. ACM, 1988.

[11] Jean Cardinal. Computational geometry column 62. *SIGACT News*, 46(4):69–78, 2015.

[12] Jean Cardinal and Udo Hoffmann. Recognition and complexity of point visibility graphs. *Discrete & Computational Geometry*, 57(1):164–178, 2017.

[13] Mark De Berg, Marc Van Kreveld, Mark Overmars, and Otfried Cheong Schwarzkopf. *Computational Geometry: Algorithms and Applications (3rd edition)*. Springer-Verlag, 2008.

[14] Pedro J. de Rezende, Cid C. de Souza, Stephan Friedrichs, Michael Hemmer, Alexander Kröller, and Davi C. Tozoni. Engineering art galleries. In Peter Sanders Lasse Kliemann, editor, *Algorithm Engineering – Selected Results and Surveys*, pages 379–417. Springer International Publishing, 2016.

[15] Satyan L. Devadoss and Joseph O'Rourke. *Discrete and Computational Geometry*. Princeton University Press, 2011.

[16] Alon Efrat and Sariel Har-Peled. Guarding galleries and terrains. *Information Processing Letters*, 100(6):238–245, 2006.

[17] Stephan Eidenbenz, Christoph Stamm, and Peter Widmayer. Inapproximability results for guarding polygons and terrains. *Algorithmica*, 31(1):79–113, 2001.

[18] Stephan Friedrichs, Michael Hemmer, James King, and Christiane Schmidt. The continuous 1.5D terrain guarding problem: Discretization, optimal solutions, and PTAS. *Journal of Computational Geometry*, 7(1):256–284, 2016.

[19] Jugal Garg, Ruta Mehta, Vijay V. Vazirani, and Sadra Yazdanbod. ETR-completeness for decision versions of multi-player (symmetric) Nash equilibria. In *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming (ICALP 2015), part 1*, volume 9134 of *Lecture Notes in Computer Science (LNCS)*, pages 554–566, 2015.

[20] James King and Erik Krohn. Terrain guarding is NP-hard. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2010)*, pages 1580–1593, 2010.

[21] David G. Kirkpatrick. An $O(\lg \lg \mathrm{OPT})$-approximation algorithm for multi-guarding galleries. *Discrete & Computational Geometry*, 53(2):327–343, 2015.

[22] Linda Kleist, Tillmann Miltzow, and Paweł Rzążewski. Is area univesality $\forall\exists\mathbb{R}$-complete? In *The European Workshop on Computational Geometry (EuroCG 2017)*, pages 181–184, 2017. Full version in preparation.

[23] Erik Krohn and Bengt J. Nilsson. Approximate guarding of monotone and rectilinear polygons. *Algorithmica*, 66(3):564–594, 2013.

[24] D. T. Lee and Arthur K. Lin. Computational complexity of art gallery problems. *IEEE Transactions on Information Theory*, 32(2):276–282, 1986.

[25] Maple 2016.1, Maplesoft, a division of Waterloo Maple Inc., Waterloo, Ontario. Maple is a trademark of Waterloo Maple Inc.

[26] Jiří Matoušek. *Lectures on Discrete Geometry*, volume 212 of *Graduate Texts in Mathematics*. Springer-Verlag New York, 2002.

[27] Jiří Matoušek. Intersection graphs of segments and $\exists\mathbb{R}$. 2014. Preprint, `https://arxiv.org/abs/1406.2636`.

[28] Colin McDiarmid and Tobias Müller. Integer realizations of disk and segment graphs. *Journal of Combinatorial Theory, Series B*, 103(1):114–143, 2013.

[29] Nicolai E Mnëv. The universality theorems on the classification problem of configuration varieties and convex polytopes varieties. In Oleg Y. Viro, editor, *Topology and geometry – Rohlin seminar*, pages 527–543. Springer-Verlag Berlin Heidelberg, 1988.

[30] Joseph O'Rourke. *Art Gallery Theorems and Algorithms*. Oxford University Press, 1987.

[31] Joseph O'Rourke. *Computational Geometry in C*. Cambridge University Press, 1998.

[32] Joseph O'Rourke. Visibility. In Jacob E. Goodman and Joseph O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 28. Chapman & Hall/CRC, second edition, 2004.

[33] Joseph O'Rourke and Kenneth Supowit. Some NP-hard polygon decomposition problems. *IEEE Transactions on Information Theory*, 29(2):181–190, 1983.

[34] Quintic function – Wikipedia, The Free Encyclopedia, 2017. [Online; accessed 14-March-2017].

[35] Jürgen Richter-Gebert and Günter M. Ziegler. Realization spaces of 4-polytopes are universal. *Bulletin of the American Mathematical Society*, 32(4):403–412, 1995.

[36] Marcus Schaefer. Complexity of some geometric and topological problems. In *Proceedings of the 17th International Symposium on Graph Drawing (GD 2009)*, volume 5849 of *Lecture Notes in Computer Science (LNCS)*, pages 334–344. Springer, 2009.

[37] Marcus Schaefer. Realizability of graphs and linkages. In János Pach, editor, *Thirty Essays on Geometric Graph Theory*, chapter 23, pages 461–482. Springer-Verlag New York, 2013.

[38] Marcus Schaefer and Daniel Štefankovič. Fixed points, Nash equilibria, and the existential theory of the reals. *Theory of Computing Systems*, 60(2):172–193, 2017.

[39] Dietmar Schuchardt and Hans-Dietrich Hecker. Two NP-hard art-gallery problems for ortho-polygons. *Mathematical Logic Quarterly*, 41:261–267, 1995.

[40] Thomas C. Shermer. Recent results in art galleries. *Proceedings of the IEEE*, 80(9):1384–1399, 1992.

[41] Yaroslav Shitov. The complexity of positive semidefinite matrix factorization. 2016. Preprint, `http://arxiv.org/abs/1606.09065`.

[42] Peter W. Shor. Stretchability of pseudolines is NP-hard. In Peter Gritzmann and Bernd Sturmfels, editors, *Applied Geometry and Discrete Mathematics: The Victor Klee Festschrift*, volume 4 of *DIMACS – Series in Discrete Mathematics and Theoretical Computer Science*, pages 531–554. American Mathematical Society and Association for Computing Machinery, 1991.

[43] Ana Paula Tomás. Guarding thin orthogonal polygons is hard. In *Fundamentals of Computation Theory*, pages 305–316. Springer, 2013.

[44] Jorge Urrutia. Art gallery and illumination problems. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, chapter 22, pages 973–1027. Elsevier, 2000.